Dust in the Wind: Challenges for Urban Aerodynamics

Jay P. Boris*
Laboratory for Computational Physics and Fluid Dynamics
U.S. Naval Research Laboratory, Washington DC. 20375-5320

Nomenclature

API	=	Application Programming Interface
CBR	=	Chemical, Biological & Radiological
CFD	=	Computational Fluid Dynamics
CFoM	=	Cumulative Figure of Merit
CT	=	Contaminant Transport
CT-Analyst	=	NRL's operational model for Contaminant Transport Analysis
DEM	=	Digital Elevation Map
Nomograf	=	NRL's new data structure for Contaminant Transport information
DNS	=	Direct Numerical Simulation
FAST3D-CT	=	NRL's Contaminant Transport LES CFD model
GUI	=	Graphical User Interface
HPC	=	High Performance Computing
LES	=	Large Eddy Simulation
MILES	=	Monotone Integrated Large Eddy Simulation
RANS	=	Reynolds Averaged Navier-Stokes
SF6	=	Sulphur Hexaflouride
WMD	=	Weapons of Mass Destruction

1. Introduction

The fluid dynamics of airflow through a city controls the transport and dispersion of airborne contaminants. It also controls environmental air quality, wind forces on buildings, and the ambient noise level due to the winds. These are urban aerodynamics problems primarily, not meteorology. The space scales are short, a few meters to at most a few kilometers. The average airflow, the dynamic fluctuations, and the building-scale turbulence are all closely coupled to the complicated geometry. Buildings create large "rooster-tail" wakes; they shed vortices dynamically and support complex recirculation zones; there are systematic fountain flows up the backs of tall buildings; and dust in the wind can move perpendicular to or even against the locally prevailing wind direction. In principle, meteorology provides the aerodynamic boundary conditions for an urban region and influences the airflow in a city but the weather can be treated as known over times of a few minutes to a fraction of an hour. Urban aerodynamics is driven by a deep, stratified urban boundary layer with significant wind fluctuations. Solar heating effects include shadows from buildings and trees and aerodynamic drag and heat exchange affected by the surface property variations and turbulent heat transport. We require time-dependent, three-dimensional Computational Fluid Dynamics (CFD) to predict accurately these unsteady, obstructed, buoyant flows and the dynamic contaminant plumes that they drive. In typical urban scenarios most particulate and gaseous contaminants behave similarly with respect to the overall transport and dispersion but the full physics of multi-group particle and droplet distributions are required for some problems.

Crisis management and reducing risks from Chemical, Biological, or Radiological (CBR) agents and pollutants released into the air from fires and accidents motivates our interest in urban aerodynamics. Computing contaminant transport (CT) accurately in cities is a difficult, complicated problem (e.g. Boris 2002; Britter and Hanna 2003). Urban landscapes have complex geometry over large areas whose widely varying temporal and spatial scales exhaust current modeling capacities. Crucial aerodynamic issues include turbulent fluid transport, interacting wakes, and boundary condition modeling. To these aerodynamic issues we should add post-processing of the CFD to distill results for practical use by responders to actual emergencies and detailed analysis of particular

* Chief Scientist and Director, Code 6400, 202-767-3055. Fellow, AIAA 2005 Fluid Dynamics Plenary Lecture

scenarios for post-event reconstruction. Computing urban aerodynamics accurately is a time-dependent, High-Performance Computing (HPC) problem.

Visualizations of Urban Contaminant Transport

Visualization serves a useful function in helping to elucidate the connections and relationships between the fluid dynamic phenomena, the geometry, and the contaminant distributions. There is certainly a subjective component to interpreting visualizations so their use can be subject to some abuse. However, things that look wrong usually are wrong, even if the converse is not always true. In this paper I will use visualization in a number of circumstances to illustrate particular issues and phenomena with rigorous scientific analyses left to other papers.

Figure 1 shows four visualizations of the dispersion resulting from an instantaneous point source release of contaminant at street level in Times Square, New York City (Patnaik, et al. 2006). This is a practical and representative example of an urban-scale CFD simulation. Buildings have been cut away in the foreground to show the evolving contaminant distribution down to street level. The figure illustrates the typical complex, unsteady vertical mixing patterns caused by building vortex shedding and recirculation, and shows the contaminated region that forms downwind in this particular ground-level release scenario. Red represents higher concentrations than orange, and orange higher than yellow, etc. The figure also shows the so-called fountain effect occurring behind a tall building. The fountain effect is the systematic migration of contaminant from ground level to the top of tall buildings on their downwind side and is one of the recognized challenges for urban aerodynamics. The contaminant fountains into the air flowing over the tops of these buildings and thus migrates downwind faster than might be expected based on the air velocity down in the urban canyons. Earlier visualizations of FAST3D-CT simulations of the flow in the Chicago downtown area (Yellen and Jacobson, 2003) discovered this phenomenon. It was also observed in experiments in Los Angeles (Rappolt 2002), and has been reported in wind tunnel studies. The fountain effect can be important because the contaminant is transported downwind much faster than might be otherwise expected during this process. This effect appears to be driven by arch vortices lying behind the buildings such as described in the well-studied problem of flow past a surface mounted cube (e.g., Castro and Robins 1977; Lakehal and Rodi 1997; Hunt et al. 1978; Martinuzzi and Tropea 1993; Snyder 1994).

Figure 2 shows a similar CFD simulation of contaminant transport for a section of downtown Houston. In these urban aerodynamics simulations the region of detailed geometry definition is always limited. This limit is shown by the flat green-brown surface extending to the horizon in Figure 2. The agent release is again instantaneous over a hemispherical volume at ground level to approximate an explosive source. During the daytime, solar heating produces an unstable atmosphere with convection currents and rolls that can spread a contaminant quickly throughout the urban "canopy," the volume above the ground but below the tops of the taller buildings. Using CFD with fluctuating winds allows the solutions to capture the voids and low-density regions that interpenetrate an evolving cloud as wind gusts and building-generated vortex transients transport uncontaminated air deep into the cloud. This natural concentration variability can be seen in both Figures 1 and 2 and is an important feature of realistic CT problems because it controls the safety factors that must be applied when considering probable exposures to dangerous airborne materials.

The Dilemma

These CFD solutions fulfill our general expectations about the fluid dynamics and turbulence. They also agree well with field trials and wind tunnel studies insofar as detailed comparisons can be made. On the other hand, using CFD technology directly in emergency assessment of industrial spills, transportation accidents, or terrorist attacks is problematic because of the very tight time requirements on operational applications. Simulations of chemical, biological, or radiological (CBR) scenarios are practical but expensive and operational users cannot afford to wait for the computations. The need for speed seems to suggest simple approximations that unfortunately produce inaccurate models. Indeed, it has been clear for a number of years that effective defense of cities, facilities, large bases, and populations against CBR incidents requires faster and more accurate prediction/assessment technology to be successful.

Figure 3 shows how many human lives can be saved if we can issue a timely evacuation warning. This warning must identify the plume centerline and councils evacuation away from the centerline across the wind. The conditions for this particular series of simulations were for a typical wind speed above the tops of the buildings, here taken as 3 m/s, and a spread angle for the plume typical of a relatively low and sprawling city like Washington DC. In this example, assume 100,000 people would receive a lethal dose if no warning were ever issued. The figure

shows that perhaps 25% of these people can avoid lethal exposure if the appropriate warning is issued in 15 minutes. 50% can escape if the warning delay is only 9 minutes. Up to 85% of these people can walk away if the warning is issued within 3 minutes. In other words 100 people per second can die for each second delayed in issuing the warning. The people should be instructed to walk away from the plume centerline perpendicular to the wind. In simple terms, a timely warning accompanied by the correct direction to walk can potentially save four out of five people who would die in the absence of a warning.

One plume prediction approximation in widespread use is based on Gaussian similarity solutions ("puffs"). This is an extended Lagrangian approximation developed for large regions of relatively flat terrain where large-scale vortex shedding from buildings, cliffs, or mountains is absent. The contaminant cloud is decomposed into a number of independent puffs with Gaussian concentration profiles and these puffs move with the local velocity and spread diffusively based on a Gaussian similarity solution with spatially constant diffusion coefficients. These solutions display only approximate aerodynamic behavior and no contaminant trapping. In the real world contaminant gets trapped in the re-circulation zones behind buildings and continues to spread laterally long after simpler models say the cloud has moved on. Additional formulae and computations, added to some of these plume models to approximate the omitted fluid dynamic effects, increase the running time. Of greater concern, these "common-use" contaminant prediction methods were never designed for situations where the input data about the source (or sources) is absent and the spatial scale is so small that problem set-up and analysis must take place in seconds to be maximally effective. Greater speed and greater accuracy are both required in today's world.

The critical problem for operational CT applications is that full-detail, urban flow simulations are currently feasible – but they are expensive and require a degree of expertise to perform. First responders and emergency managers coping with contaminant release threats cannot afford to wait while actual simulations and data post-processing are carried out either on site or elsewhere. Thus, the dilemma has been to choose either a fast model giving questionable results in important circumstances or to wait a dangerously long time for more accurate results. Resolving this dilemma is an important challenge to urban aerodynamics. This paper discusses the aerodynamics and physics of this problem and presents a workable answer to this dilemma. Using new fluid-dynamic principles, an urban-oriented emergency assessment system called CT-Analyst® has been invented to produce accurate results for airborne contaminant scenarios nearly instantly, e.g. (Boris, et al. 2002, 2005). Designed to predict airborne contaminant flow including Chemical, Biological, and Radiological (CBR) threats, CT-Analyst has some unique capabilities and gives HPC accuracy while running much faster than other current alternatives. Section 5, entitled ***Theory and Practice,*** explains how applied aerodynamics enabled this new capability.

Why Use Computational Fluid Dynamics?

Given the computing burden usually associated with CFD, why should we use it to compute urban aerodynamics and contaminant transport? There are a number of different approximations, models, and approaches that might be considered and CFD is just one of them. Sections 2 and 3, ***The Dynamics of Contaminant Transport*** and ***Aerodynamics and Meteorology***, answer this question from several perspectives. The simple answers given here set a framework for those discussions.

CFD is the approach to use when adequate time and computing resources are available. Scientists with the required skills and resources should set high standards by doing as good a job as possible to solve urban aerodynamics problems accurately. The fewer the possible sources of error, the more reliable the results will be. The physics and fluid dynamics of the real problems are time-dependent so the starting point should be to resolve these effects and CFD does this. Using the CFD approach, there is some hope of pushing the frontiers of science back by staying close to fundamental principles. In this case this means trying to solve the high Reynolds-number limit of the Navier-Stokes equations with as few limiting assumptions as possible.

In the name of speed and simplicity, sponsors and operational personnel often prefer running models on laptop computers or small workstations. This focus also encourages intellectual creativity in developing and testing new models, new phenomenologies, and simple approximations. Such efforts can be more satisfying than slaving away with a large computer to get the detailed geometry, fluid dynamics, and physics correct. These latter efforts are absolutely required, however, to constrain intuition and other flights of fancy by the physics and fluid dynamics of the problems under attack. Many simplifications reduce the computer requirements and lead to faster solutions but the quest for reliable solution accuracy means that one should not average over features that can be resolved in space and time. Mathematical or physical models that incorporate such averaging are automatically limited in accuracy and thus ultimately limited in their ability to make reliable predictions. Computing power is always

increasing so it makes sense to use an approach and technology that permits progressively better answers when appreciably more computing is available. Complex geometry CFD with a large-eddy subgrid turbulence treatment meets these resolution and physical convergence requirements.

Contaminant transport and dispersion are not local phenomena so the actual geometry is crucial. Where a contaminant cloud goes and how long it stays there is a nonlinear, time-dependent integration procedure over the entire volume of interest. This integral involves fluid dynamics in regions that the contaminant never visits because these regions can affect the velocities and turbulence in the regions that do become contaminated. This means, in particular, that a volume representation of the turbulence with the correct geometry taken into account is needed because turbulent transport dominates the dispersion on all reasonable scales. Turbulence modeling only on the surface of bodies is inadequate.

2. The Dynamics of Contaminant Transport

The fluid dynamics of airflow and contaminant transport in urban regions is an intrinsically unsteady flow problem in complicated geometry. In this section the important physics and fluid dynamic effects are discussed and related to requirements for accurate modeling. As a way into this discussion, I list several observations that provide insight into simplifications that help either the modeling problem or help us focus more efficiently on the key issues. These are taken from studying hundreds of detailed simulations, from discussion with first responders and emergency officials, and from discussions with experimentalists and colleagues. These simplifications take some difficulties out of play and help make the contaminant transport problems in urban aerodynamics computable in a practical way.

Simplifications

Neutrally buoyant tracer gases give the largest downwind doses because they don't tend to rise above the buildings due to additional buoyancy (as in fires) and they don't tend to fall to the ground. Thus there is more neutrally buoyant tracer in the air downwind, as a fraction of the amount released, than with any other kind of contaminant. Furthermore, it is simplest to assume that the contaminants are not absorbed, they don't react with their surroundings, and they don't deposit on surfaces. This takes a lot of information about material properties such as porosity out of play. Because neutrally buoyant tracers lead to the largest inhaled doses downwind, this most conservative case is also the easiest to simulate. The inhaled dose from other types of threats will be less severe because the agent's decay or the particle settling reduces the amount in the air to be inhaled. An exception can occur when a large amount of a cold gas is released quickly. Then the density of the contaminant is large enough to affect the buoyancy and the plume may form a relatively thin layer of high concentration near the ground. This effect dissipates, however, with mixing and dilution of the contaminant downwind.

Buildings control the lateral spreading of contaminant to a great degree. The examples shown in Figures 1 and 2 are characterized by migration of the contaminant across the prevailing wind in the recirculation zones of closely spaced buildings. This means that the contaminant plume envelopes are controlled by and locked to the static geometry. Variations of the relative strength of wind gusts, for example, or differences between day and night dispersion in urban areas leads to only small (a few percent) variations in the lateral extent of clouds. The contaminant arrival times, concentrations at any given time, and decay rates show greater sensitivity to these environmental effects but the keep-out region, the first concern of emergency managers, is little affected. In fact, the dependence on the building geometry is so important that very different plume extents can result from small displacements of the source location, as shown in Figure 10 and discussed in Section 4, ***Challenges for Urban Aerodynamics***.

There is very good vertical mixing in cities. The examples of Figures 1 and 2 show this and it has been observed in early urban aerodynamics calculations (Boris, et al, 2001b). Good vertical mixing was one of the first highlighted results of the Urban 2000 field trials conducted in Salt Lake City (Allwine, et al. 2002). This behavior means that very little three dimensionality needs to be reported out for emergency uses even though the computation of the underlying dynamic flow field must be three dimensional to properly account for the lateral and downwind spread rates.

Dispersing plumes have relatively sharp edges in the early phases because the major transport mechanism is convection in the large eddies and coherent flows channeled between, around, and over buildings. Diffusion in the strict sense is not a controlling factor so the edges are not so quickly blurred. The consequence of this is that large

errors in the amount and even type of contaminant present have little effect on the area that will have to be declared as a contaminated region. The computed plume envelopes are only a little larger when the definition of the plume edge is based on a hundred-fold increase in the amount of contaminant.

Since a prime motivator for our interest is computing CT in urban regions to support decisions during a crisis, the fact that a near optimal response actually requires very little information about the source is very important. Where the source is located is important because it controls the arrival time, shape, and duration of the contaminant concentration curve but the amount and even the contaminant itself is of less significance. To minimize the amount of contaminant a person inhales is an achievable and meaningful objective. Adopting this as a goal ensures that the amount of an agent released is of little importance to optimal defensive actions in the short run. If the actual amount turns out to be ten times what we thought it was, the strategy to minimize the inhaled dose won't change at all. The inhaled dose will still have been minimized even though the actual amount turns out to be ten times larger. This also means that the actual material is of correspondingly little consequence to what must be done in the first few minutes.

The FAST3D-CT CFD Simulation Model

Direct numerical simulation (DNS) is defined as the time-dependent solution of the full Navier-Stokes equations down to the small scales where turbulence energy is being dissipated by the molecular viscosity. DNS is out of the question with today's computer technology for simulations of urban aerodynamics. The standard industrial approaches, such as Reynolds Averaged Navier-Stokes (RANS), occupy the other end of the CFD fidelity spectrum. These models simulate the mean flow and approximately model the effects of turbulent scales (Hendricks et al. 2004). Such models may be unacceptable for urban CT modeling because they are unable to capture the inherently unsteady plume dynamics driven by the urban geometry. Large Eddy Simulation (LES) constitutes an effective intermediate approach between DNS and the RANS methods (Sagaut 2004). LES is capable of simulating flow features that cannot be seen with time-averaged (steady-state) methods, such as significant flow unsteadiness and localized vortex shedding, and provides higher accuracy than the industrial methods at a lower cost than DNS. LES solutions converge to the solutions of the Navier-Stokes equations as resolution is increased, whereas RANS results generally do not. Because the larger-scale unsteady features of the flow govern the unsteady plume dynamics in urban geometries, the LES approximation is arguably the proper class of models to adopt.

The NRL FAST3D-CT model provides a clear demonstration that detailed CFD, in the LES approximation, can be executed with finite resources for urban aerodynamics at the building and street scale. Early descriptions of the capability appear in Stefaniw, et al. (1998) and Boris, et al. (1999). The fluid dynamics and turbulence models in FAST3D-CT all use the nonlinear, monotonicity-preserving Flux-Corrected Transport (FCT) convection algorithm. Boris (1971), Boris and Book (1973, 1976), Boris et al. (1993), Young, et al. (1993), and Patnaik, et al. (2003) describe FCT and the treatment of complex geometry underlying FAST3D-CT. Oran and Boris (2001) consider many of the auxiliary algorithms and coupling procedures in detail and Kuzmin, et al. (2005) survey and analyze world-wide progress in FCT algorithms.

The Monotone Integrated LES (MILES) turbulence model (Boris 1989; Boris, et al. 1992; Oran and Boris 1993) has also been adopted for FAST3D-CT. Because of its high computational efficiency, MILES is well suited to CFD-based plume simulation for urban-scale scenarios, an application where classical LES methods are expensive. See Grinstein and Fureby (2004) for a recent review of MILES and related Implicit LES models and the book ***Implicit Large Eddy Simulation: Computing Turbulent Flow Dynamics*** (Grinstein, Margolin, and Rider (eds) 2006) for detailed discussions and validation studies of this MILES approach.

FAST3D-CT is the time-accurate, high-resolution CFD model that underpins the CT-Analyst operational model described below. Boris (2002) presented FAST3D-CT in some detail and introduced its use as the foundation for CT-Analyst. This article and the chapters by the NRL team (Patnaik, et al. 2005; Patnaik, et al. 2006) provide information about FAST3D-CT use for urban aerodynamics in greater detail than will be treated here. FAST3D-CT is based on a scalable (parallel processing), low dissipation, 4^{th} order phase-accurate Flux-Corrected Transport (FCT) convection algorithm (Patnaik, et al. 2005). Buoyancy is treated by a simple potential-temperature formulation and the inflow temperature profile can be established by several input data parameters. Solar heating is a factor. A ray-trace algorithm allows buildings and trees to cast shadows on each other with heating differentiated by the surface type, the surface's orientation relative to the specified sun direction, and whether the sun is being partially or fully blocked. A primary difficulty is the effective calibration and validation of these physics models since much of the input needed from field measurements of these processes is typically sparse or nonexistent.

Further, even though the individual models can all be validated separately, the larger problem of validating the overall simulation code has to be tackled as well.

Figure 4 shows a top view of the contaminant concentration for a 2 km by 2 km area of downtown Chicago resolved with 6-meter finite volume cells. The concentration is plotted 10 feet off the ground and the concentration values are color coded with yellow being more concentrated than green, followed by light blue and dark blue. This horizontal cross-section, shown 18 minutes after the source was released in a 3 m/s wind from the 320^0, is the lowest of 55 planes from a fully three-dimensional FAST3D-CT simulation. This CFD result corresponds in detail with the CT-Analyst example shown as Figure 13 and discussed in the Section 5 below. This figure illustrates the complexity and resolution of the flow geometry (see also Figures 1, 2 and 10). A typical run with the FAST3D-CT model for a complex urban area of 30 square km resolved with 6 m cells takes half a day on a 20-processor SGI Altix. This is significantly faster than classical CFD models due to the savings achieved by MILES as well as other algorithmic improvements.

In FAST3D-CT the wind boundary conditions are allowed to fluctuate. A deterministic model for a synthetic realization of these input fluctuations is included with two free parameters, a spatial-scaling factor and an amplitude-scaling factor. These are specified as input parameters, allowing a user to approximate different types of atmospheric conditions. The fluctuating wind model we have used is an analytic function defined throughout the computational domain. It serves as initial conditions at the beginning of a run and is updated thereafter at the side and top boundaries of the domain. All fluctuations occurring within the domain during a simulation are naturally sustained by the vortices shed from buildings and the nonlinear evolution of the impressed fluctuations at the inflow boundaries as it is convected through the domain.

Three types of essentially incompressible motion are superimposed at several different wavelengths to construct the analytic wind-fluctuation function. A coherent shearing motion with a sinusoidal structure typical of meanders is impressed as the first type of motion transverse to the average wind direction. Superimposed on this is a set of horizontal pancake vortices at several scales to represent the types of flows observed in stratified fluids. The third motion is longitudinal roles with finite vertical and horizontal extent to represent the rollers found in typical boundary layers. These rollers are stretched by the average wind shear but no attempt has been made to correlate the location of these rollers with the details of the topography. This correlation, if it occurs, arises naturally. The multiple scales for each of these motion types were made incommensurate, permitting nonlinear interactions to guarantee an overall chaotic boundary condition with stable statistics. In addition, a nonlinear term was included in all sinusoidal dependences to force a broad spectrum of impressed fluctuations.

The vertical dependence of these resolved-scale fluctuations is a superposition of two functions, one for the unobstructed flow and one to provide additional fluctuations due to the buildings. The unobstructed component is largest in the center of the domain away from the ground and below the atmospheric boundary layer. The building component is largest near the tops of the buildings whose general disposition and height determines the shape of the average inflow urban boundary layer. When this composite turbulent field is allowed to evolve by flowing over 500 to 1000 meters of actual city geometry, initial errors and inconsistencies give way to a more accurate and self-consistent flow. This has also been observed directly in wind tunnel validation studies (Patnaik, et al. 2005, 2006).

In FAST3D-CT the inflow-outflow algorithm is applied over the entire boundary and combines the average and fluctuating components of the wind. This condition morphs continuously from an approximate characteristic treatment for inflow to a simple constrained extrapolation for outflow. At side boundaries the analytic boundary values are used to gently “steer” (relax) the self-consistent values from inside the domain toward conformance with the scales and conditions being suggested at the boundary. As the flow direction is continually changing and even reverses in some places, the boundary condition adapts from one to another of these conditions continuously. We intended this model to provide an adjustable framework for incorporating progressively more information about the actual wind conditions for each urban area being considered and to act as a model for performing generic sensitivity studies (Patnaik, et al. 2003). When suitable data becomes available, this model should be the subject of an extensive calibration study to convert it to a trusted empirical model capable of being tailored to each urban application.

The FAST3D-CT model also has algorithms and input for important droplet and aerosol physics including evaporation and conservative re-lofting off the ground - when the particular problem being solved warrants including these effects. The model can be run in a single-phase (vapor) mode, in a two-phase (vapor and particles or vapor and droplets) mode, or in a three-phase (vapor, particles, and droplets) mode. FAST3D-CT is a "multi-group, multiphase" model in these latter applications. Each scenario can have many separate species and they can react

chemically - since FAST3D-CT started as a reactive flow code. Each species can, in principle, contribute to the density with buoyancy effects being taken into account though this is not the normal use of the model. The equations being solved can therefore include pyroclastic flows such as occur in volcanic releases or when the World Trade Towers fell.

An important issue remaining is what numbers to input to describe all these physical effects. Though the physics itself has been included rather simply, many of the input values are tenuous or even unknown for these processes, particularly in the case of unknown contaminant materials. Therefore, simulations based on plausible worst case threats are used to provide the compressed databases that drive the CT-Analyst operational tools described below. Meanwhile, the detailed CFD simulations can be applied directly to sensor system optimization, to computations supporting the defense and design of buildings, to physics and environmental sensitivity studies, to forensics, and as a source of virtual field trials for microscale and nanoscale atmospheric fluid dynamics and aerosol physics relevant to urban design and airborne contaminant defense.

Uncertainty, Variability, and Fluctuations

"Uncertainty" describes the situation when some variables are simply unknown. "Variability" applies when the particular values assigned to a quantity at different times and places appear random but these values are being drawn from a known statistical distribution. What agent will be used in a particular attack is uncertain, as is the location and time of the release. On the other hand, the atmospheric turbulence and the gusts in the wind can be measured, in practice and in principle, in almost any environment. Therefore the air velocity and wind gusts are variable but not uncertain. In practice the wind variability in an urban setting is seldom measured accurately enough to be the basis for validating CFD models but appreciable information is or can be made available.

Even if the wind variability were very well characterized, however, it does not follow that the corresponding variability in contaminant concentrations measured at some location far from the source can be known. The downwind concentration is determined by integrating a complex, nonlinear continuity equation that uses the complete time- and space-dependence of the air velocity. The concentration value at a point, therefore, has many possible values depending on the particular realization of the airflow time history that is chosen from the ensemble of possible urban aerodynamics solutions corresponding to the current environmental conditions. Often field trials emphasize long duration releases and long duration sampling intervals to reduce measurement variability. Many accident and attack scenarios, however, involve localized short duration releases sometimes called "acute" events. Tracer ES&T, Inc. performed a series of short duration "acute" SF6 releases in downtown Los Angeles in 2001 (Rappolt, et al. 2002). This region and data set are the subjects of a validation study discussed in Section 5.

Here, I introduce a couple of figures from this study to show that the differences from one realization to another are substantial – even when the releases are five minutes long. Inspect the eight panels in Figure 5. This figure is for a case with moderate wind fluctuations and the wind from 170^0 with a speed of 3 m/s. The conditions for the computations in each panel differ only in the time at which the source was released. The wind fluctuations and vortex-shedding phase differences between realizations account for all the differences in the observed plume shapes and concentration distributions. Contaminated areas can differ by factors as large as 40-50% as seen by comparing Source 2 and Source 8. We rightly conclude that there is no single true experimental answer to be compared with simulations or models. Measuring multiple realizations of the experimental conditions is usually not possible at city scale, leaving us with no quantitative yardstick to compare experimental and computed concentration values (Obercampf and Helton 2002; Oberkampf and Barone 2004).

Figure 6 was computed by turning the wind fluctuation amplitude factor in FAST3D-CT but keeping all other factors the same as for Figure 5. This turns wind gusts off all together but brings only a small reduction in this naturally occurring variability observed because building vortex shedding still provides appreciable turbulence and the neutral boundary layer is unstable from solar heating. One obvious result of including the fluctuations is to hasten the release of contaminant from building recirculation zones. This effect is visible by comparing the corresponding plumes in Figures 5 and 6. There is a little more orange near the source location in the panels of Figure 6, indicating slightly higher concentrations there. This effect would be even stronger if the solar heating had been turned off by computing a night-time condition for Figure 6. This naturally occurring variability is a major challenge for urban aerodynamics and so is discussed in Section 4 below and again in Section 6.

So What's the Big Deal?

Computing urban aerodynamics and the contaminant transport that results is not an impossible task. It is easy enough to compute the detailed flow over large areas of a city with reasonable resolution that oversimplifying either the geometry or the physics isn't universally necessary. The FAST3D-CT model, for example, requires an hour of computing per source to compute a one-hour scenario when running on 16 processors. As a consequence, there are a number of teams today performing urban aerodynamics computations using different models and approaches. Camelli, et al. (2004, 2005) describe a finite element model embodied in a code called FEFLO-URBAN. Aliabadi, et al. (2004) and Baffour, et al. (2004) describe a time-implicit finite element model that has been run both time accurately and with very long time steps to considerably reduce the computing time.

The FEM3MP model developed at Livermore can be run in both RANS and LES modes (Chan, et al. 2000; Calhoun, et al. 2004) and extensive comparisons of urban aerodynamics and CT have been performed to study the differences. Koomullil and Soni (2001) have performed well-resolved, time-dependent simulations of contaminant transport for New Orleans. FLACS (Hanna, et al. 2004) is another RANS model being run in time accurate mode. A distributed porosity algorithm is used to implement the complex building geometry in FLACS, not one of the usual aerodynamic approaches to represent solid bodies. This model has been used with fluctuating wind boundary conditions intended to simulate meanders in the wind and thus can also consider the concentration variability problem.

There is an even larger number of tested aerodynamics models that could be adapted to perform similar urban computations including hybrid RANS/LES models and so-called Detached Eddy Simulations. Atmospheric models have been adapted for complex geometry, such as HIGRAD (Brown, et al. 2001; Smith, et al. 2002), and commercially available CFD models, such as FLUENT (Huber, et al. 2003), are also being used for some urban computations. I apologize that it isn't possible to be exhaustive in listing all the models and efforts being applied to this problem but from the perspective of this paper, RANS approximations, other steady-state models, and various reduced-geometry and/or reduced-physics models really aren't necessary as components of a primary "predictive" capability. We should be aiming higher. In one manner or another, these models average over large-scale coherent eddies in the flow that can actually be resolved. Because steady-state models, particularly in the finite element formulation, seem to be more costly to run than finite-volume CFD models, there seems to be no corresponding performance gains for giving up the physics of time-dependent, building-scale turbulence. A couple of thousand iterations are usually required to find a steady-state flow field and each of these iterations is more expensive than a time step in an explicit LES model. These averaging models introduce additional parameters corresponding to the phenomenological models being incorporated to replace the missing physics. For example, using a time step of 30 seconds, as allowed by an implicit time integration algorithm, suppresses all phenomena that occur at this scale and faster. This reduces the dispersion substantially. Decreasing the resolution at the same time introduces a lot of numerical diffusion that spreads the contaminant more broadly. In some cases these two errors can appear to cancel but this is not a viable predictive solution. Doing complex-geometry flow solutions in a time implicit or steady state CFD framework seems an expensive way to provide a framework for fluid dynamic phenomenology.

3. Aerodynamics and Meteorology

Solving problems in urban airflow is about three quarters aerodynamics and about one quarter meteorology. On the scale of city streets, buildings, or trees, the problem is primarily aerodynamics. Vertical surfaces and edges become vorticity sources in complex geometry. Even buildings of modest height are higher than the boundary layer over open ground so vorticity and turbulence is being injected directly into the free flow occurring many cells above the ground plane of a meteorological model. When the buildings are tall, it can take kilometers of distance downwind for vorticity, generated by a ground plane roughness model intended to introduce building effects, to reach the height in the computational grid that it should have at its creation. Such ground conforming meteorology models and "urban canopy" models are simply under-resolved phenomenologies when it comes to reliably computing urban aerodynamics.

In this context, body fitted coordinates and finite-element representations of the geometry are also of little use for two reasons. First, urban surfaces are usually irregular and rough so laminar and turbulent boundary layers generally become separated at corners and sharp edges. Computing this separation and the fluid dynamic circulation that results follows from the simple conservation laws. Detailed treatment of the surfaces and edges is unwarranted. In any case, the geometric details that might be resolved in principle are usually quite unknown. The second reason is that turbulence throughout the volume is very important. The dispersion of a contaminant is governed by what

happens away from the obstacles. The more open the terrain becomes the more this is true. Many of the usual aerodynamic applications benefit from considering the turbulence as associated with surfaces but this CT application does not. In particular, concentrating numerical resolution near the building surfaces buys very little and increases the numerical costs. In fact, since concentrating numerical resolution near buildings means reduced resolution in the volume for the same length runs, the variable resolution appears to be a net loss.

Meteorology gives the boundary conditions for specific urban aerodynamics computations. These conditions must be applied at all the boundaries throughout the duration of a run, not just as initial conditions at the upwind boundary. The sides, top, and outflow edges of the system must all be treated in a uniform manner because strong wind fluctuations can temporarily change an inflow boundary to an outflow boundary. In developing and using dynamic, time-varying boundary conditions, there is still a lot of computational art as well as science required.

Times of interest are only typically a few tens of minutes to an hour or two, particularly in operational applications of emergency and planning tools derived from the detailed CFD. Therefore we should measure the environmental conditions *in situ* rather than trying to predict them ahead of time. Meteorology and atmospheric prediction has an important role in urban aerodynamics, producing the statistical distributions and wind fluctuation morphology throughout the region that can be distilled into synthetic boundary conditions for the urban aerodynamic models. This does not call for direct coupling of the models but instead spawns a major simulation program founded in the development and production running of microscale atmospheric models.

Urban aerodynamics is really nanoscale atmospheric modeling with grid resolutions of a few meters up to maybe 10 meters. Mesoscale models today run normally at 100 times this scale. Some claim accuracy down to 300 meters or so but this claim is clouded by their solution of inconsistent equation sets, by poor treatment of the geometry, by the adoption of urban canopy phenomenologies, and by physics models that average over resolvable dynamic phenomena.

Therefore, there is a pressing need for microscale meteorological models with horizontal resolution in the 30–50 meter regime that can be run with typical boundary conditions from mesoscale models. These microscale simulation codes should be used to produce the morphology and appropriate statistical distributions needed to calibrate and validate synthetic, dynamic boundary condition models such as that implemented in FAST3D-CT. Bits and snatches of data and single runs are of little or no use. Turbulent kinetic energy strengths, expressed as a function of height, are also virtually useless until accompanied by correlation lengths and characteristic frequency spectra.

Directly nesting an urban aerodynamics simulation into boundary conditions given by a single weather-model run, even assuming the weather model produces boundary conditions in a useable dynamic form, delays the results and reduces the generality. This nested-model approach is entirely appropriate for forensic studies, however, when we need to understand exactly what did happen in one particular event. An extensive modeling and field test program is needed. In a city the morphology of the wind fluctuations usually depends on the wind direction. In Chicago, the wind off the lake has an entirely different behavior than the wind coming in over the plains and the suburbs (Iselin, et al. 2002). The new micrometeorological models should be able to synthesize the information needed to drive the nanoscale urban aerodynamics runs from all directions.

The suggested production program is not forecasting in the usual sense and thus falls outside the worldview of most computational meteorologists. This effort would be more in the nature of aerodynamic parameter surveys to provide a design database. We need to know limiting cases and typical fluctuation structures in different weather conditions, at different times of the day, and in different seasons. At the same time an extensive experimental program should be started to collect time- and space-resolved wind data to validate the models and to provide direct empirical measures of the required data for key regions and cities.

We have a mechanism, the dispersion nomografs underpinning CT-Analyst and described in some detail below, to take full advantage of this new micrometeorological information and these new capabilities. This new approach to real-time emergency response requires that short and moderate wavelength wind gusts be impressed at an upwind boundary. These gusts play a strong role in the decay of contaminant, as discussed and illustrated by Figure 8. Thus, they ought to be included even though they are less important than the vortices and turbulence being shed from buildings in controlling the shapes of contaminant plumes. Please inspect Figures 5 and 6 again. The important space scales (tens to hundreds of meters) and time scales (seconds to minutes) can be resolved easily by CFD models that are resolving the buildings and should not be averaged or approximated by simpler models of steady state approximations.

4. Challenges for Urban Aerodynamics

Understanding and predicting urban airflow are relatively new applications for aerodynamics and they add new challenges to the list of issues that arise from aircraft aerodynamics such as separation, turbulence, and aero-acoustics. In Section 2 above, useful, problem-specific simplifications were discussed that make the job tractable and that illuminate a path to fast, operational solutions. It is now time to consider the challenges.

Fluid Dynamic Challenges

Urban aerodynamics problems are intrinsically unsteady and turbulence acts extensively throughout the volume rather than primarily at surfaces. The turbulence that is generated at surfaces arises from protrusions, ridges and sharp edges and is not the natural fluid-dynamic evolution of laminar boundary layers on smooth surfaces. The geometry determines where flow separation will occur. Therefore dispersion of a contaminant in these flows is dominated by convection in coherent vortex structures that can usually be resolved in computing the free flow. Shed vortices are carried by and help define the wind, and they continue the dispersion of contaminants far from the ground and buildings. This puts a premium on numerical methods that preserve unsteady free-stream turbulence without strong decay from nonphysical numerical artifacts. Under these conditions, it isn't wise to expect many of the classical methods to work well.

Variability and Concentration Variance

The intrinsic unsteadiness of urban flows creates other challenges. Because the inflow conditions are variable and the shedding of vortices behind buildings adds to that variability, a particular tracer-release experiment cannot be exactly duplicated. Even assuming that the statistical fluctuation spectra for the prevailing winds are known, the particular realizations will all be different, as seen in Figures 5 and 6. In the real world the separation between what can be treated as fluctuations drawn from a known spectrum and what variations should be treated as a change in the spectrum is not made easily. This is just a real-world consequence of the separation-of-scales problem that also arises in defining RANS models. In FAST3D-CT we input the 10-minute average winds and treat them as given steady profiles while adding the specific, dynamic realization of space- and time-varying fluctuations on top of these profiles.

In principle, the wind fluctuation spectra and the spatial correlations can be measured but the variability in contaminant concentration values at a location well downwind of a source really cannot. This concentration variance, however, is the key factor that controls the expected uncertainty in measurements and predictions. It takes a few minutes for distinct experimental releases at a single urban location to de-correlate because the low frequency components of the velocity fluctuations extend over several minutes. This means that only two or three independent realizations can be initiated before the background wind and weather conditions themselves are changing systematically. Two or three experimental trials, however, do not give enough data for good statistics. This challenge to determine the concentration variability further taxes simulation technology to find an answer because the information cannot be extracted directly from even perfectly measured wind fluctuations. Further, because of this missing information, we can't even compare simulations and field trial data quantitatively on a point-by-point basis.

A detailed numerical integration is required to determine this concentration variability computationally, even when the wind variances and governing structure are known perfectly. A full simulation with all the geometric factors can determine the time-varying winds numerically everywhere the contaminant goes, but it must be followed by an integration of the contaminant transport equation. Computing multiple realizations of a single release can approximate the concentration variance - for a computational price. Using a model this way can give important insights and even inspire confidence in the computed results. Using this simulated concentration variance as a yardstick to validate a CFD model against the limited field trial data, however, approaches being a circular argument.

Figure 7 illustrates what the issues are here. The four yellow areas in the figure outline four different realizations of a possible release at some fixed time after release, say 10 minutes to be concrete. The release point in the city is the same for each realization but the plumes have been separated in space for illustration purposes. The wind speed and direction is the same in each case but the details of the fluctuations are different for each member of the ensemble. These four plumes differ in shape, as in Figures 5 and 6, and these differences are related to the concentration variance yet to be determined. Each realization is shown superimposed on a single "plume envelope,'

the union of the four individual realizations. Each of the four realizations shown fills a different fraction of the plume envelope, depending on the particular history of fluctuations that defined that realization. The plume envelope is the prediction of the contaminated region emergency managers would like to have in an operational context because the likelihood of contamination outside that region in the first 10 minutes is very small.

If any one of the possible plume realizations, whether computed or experimental, is used for this operational purpose, there is a good chance that contamination will occur in places it was not predicted. Therefore, being safe and conservative means overestimating the contaminated region from a particular event to account for the known variability in exactly where the contaminant will go. It is clear that minimizing the amount of this overestimation is a design goal and a Cumulative Figure of Merit (CFoM) has been defined to do this (Boris, et al. 2005; Young, et al. 2004). This CFoM is defined to be unity only when the union of the plumes has at least one realization with a local concentration value above threshold for each point of the plume envelope and when no concentration values above threshold are observed to lie outside the plume envelope. The plume envelope is a function of time and so the CFoM will also be a function of time. The next two sections return to these issues from the perspective of a fast running operational model (Section 5) and validating urban aerodynamics computations (Section 6).

Improved Fluctuating Wind Model

I described the FAST3D-CT wind fluctuation model briefly in Section 2. That discussion and the importance of understanding variability in the wind and the contaminant concentrations, points to another challenge for urban aerodynamics, constructing and validating a fast, more-physically motivated model for computing fluctuating wind realizations for use as LES boundary conditions. Any improvements to the current model should be based on data we currently do not have in hand. Doing this job properly is a several year research effort and would benefit greatly from an experimental program measuring the time-dependent behavior and structure of the wind fluctuations. The actual morphology (structure) of the fluctuations has to be understood first and then described in a manner that can be parameterized simply and computed efficiently. Any wind fluctuation model that does not take some account of the structure of the terrain and buildings upwind of the CFD computational domain will be inadequate because these clearly play a significant role.

Decrease of Contaminant Densities with Time (Decay)

In the unsteady, variable environment of airflow over a city, contaminant decay is a complicated process and provides another difficult challenge. Contaminant trapping and release from recirculation zones is controlled by the geometry, by the natural wind variability, and by the solar heating. We would like to be able to predict with some reliability how high a pollution level is to be expected and how long a polluted area will stay polluted after the source has been removed. Current estimates vary greatly (Pullen, et al. 2005; Gaydos and Nyden 2004). Figure 8, computed by FAST3D-CT, shows how quickly a contaminant is swept out of a city by the dynamic airflow. The vertical axis is the decay time in minutes, describing how quickly the contaminant is scoured out of the urban region. The decay time can vary by a factor of four or more due to solar heating variations from day to night and from differences in the average wind fluctuation strength. The horizontal axis of the figure is the relative strength of the wind fluctuations imposed at the boundaries of the domain being computed. On the left the fluctuations are about 20% of the steady component of the wind and on the right the fluctuations are as large as the average wind.

For each of six different “environmental conditions”, twelve ground-level sources were released with four realizations of sources released at each three different locations in the downtown area of a city. These source locations and the scale lengths of the wind fluctuations were held fixed for the six different runs but the fluctuation strengths and the solar heating (day and night) were varied as indicated in the inset table. The value of the exponential decay time is plotted for each realization and source as a diamond-shaped symbol whose color corresponds to one of the six environmental conditions. The three source locations for each condition are distinguished by a slight horizontal offset of the ensemble of four decay times, connected by a vertical line in each case. Also plotted with each ensemble is the average, shown as a filled circle.

The dark blue diamonds (40% fluctuations at night) should be compared with the light blue diamonds (40% fluctuations during the day). Similarly the purple diamonds (60% fluctuations at night) should be compared with the unfilled red diamonds (60% fluctuations during the day). The results show that the contaminant decays two or three times faster during the day than at night under otherwise identical conditions. One can also see that the decay times get systematically shorter as the wind fluctuation amplitude is increased (from left to right in Figure 8). Increase the wind fluctuation strength imposed at the boundaries by a factor of five, however, decreases the decay time by at

most a factor of two. This trend is emphasized by the light blue shaded bar through the center of the four daytime data sets in the lower portion of the figure.

These results reinforce the great importance of concentration variability. The spread in the decay time for nominally identical realizations is quite large, often a factor of two or more. Slight changes in the time when a contaminant is released may allow a wind gust to kick the bulk of the contaminant above the buildings in one realization and into a protected courtyard in another.

Contaminant cleansing is an important challenge. It is possible, in principle, to structure buildings so that trapping regions and recirculation zones are quickly flushed and the air replaced by outside air. This could reduce the exposure of populations in CT situations and the overall quality of the air in ordinary everyday life. The goal is to reduce the decay times during the day and the night by creating channels of fresh air into the recirculation zones to augment the slow exchange attending the natural vortex shedding. Modifying the geometry or relative placement of some buildings is one possible way to decrease the decay time of trapped contaminants and thus to reduce the risks from airborne contaminants while decreasing the adverse effects of pollution. For example, a channel could be provided at ground level from the upwind side of a building, for example through a courtyard, to the downwind side. This channel would allow air that normally would stagnate for a while in front of the building to flush the courtyard and the recirculation zone behind the building, reducing the decay time appreciably.

Wind channel deflection can also be accomplished in the building design phase. Designs could include methodology to move wind gusts off the street level. This trades off against cleansing the air but wind deflection would be reducing the unpleasant effects of fast hot winds in the summer and fast cold winds in the winter.

Understanding Urban and Building Effects

Buildings take up volume in a city and can also take up some of a released contaminant. In Washington DC about 28% of the area is covered by construction. A much smaller fraction of the volume below 100 meters is filled. In Manhattan these fractions are much larger. An appreciable fraction of a contaminant released in a city can infiltrate buildings when they are large and close together. Further, the density of the contaminant remaining in the streets is higher because the dispersion is being hampered by the buildings, relative to spreading in an open area. As the contamination in the streets moves on, the appreciable amount captured in the buildings slowly leaks out, maintaining the concentration in the streets at a higher level than expected for a long time. These effects have never really been studied and are not accounted for in today's common use models.

Traffic, air conditioners, heating system exhausts, subways, and opening and closing doors all introduce small perturbations to the airflow and to the temperature of the air. The heat being released in buildings, or their cooling during the summer, ensures that their outer surfaces will be at a different temperature from the air passing by. These modifications of the flow may be large locally and they certainly can add up over time and distance. If this "urban heat island" is considered at all in meteorological models, it is treated through area-averaged coefficients with a scale in kilometers. In urban aerodynamics this information should be differentiated, at least in principle, down to the street scale. In a large-scale simulation covering kilometers some averaging will be used but it will not be the same as used in the meteorological models. For example, a tall building in winter heats the passing air along the full extent of the sides of the building. However, if only the footprint area is used, the heating would be much less and would occur at the ground plane.

Another challenge is the numerical treatment of traffic. Traffic can act as an appreciable turbulence generator even when the wind is not very slow. An important challenge is to understand and quantify these various turbulence effects. Added turbulence at street level can accelerate the rate at which low-lying contaminant percolates up into the air near the building tops where systematic flow carries it away faster than if it were to remain at street level.

The Fountain Effect

Visualizations of FAST3D-CT simulations of the Chicago downtown area such as reproduced in Figure 9 (Yellen and Jacobson 2003) showed an intriguing effect, the systematic migration of contaminant from ground level up the backs of tall buildings followed by continuous ejection into the air flowing over the tops of these buildings. In visualizations of the contaminant concentration this looks just like a fountain, as also shown in Figure 1, and we call it the "fountain effect." It is important because the contaminant can be transported downwind much faster than might be otherwise expected during this process. Similar observations have been made in experiments in Los Angeles (Rappolt 2002 and private communication) and have been reported in wind tunnel studies. It is a real

challenge to urban aerodynamics to explain the mechanism behind this effect.

What appears to be happening is this: the vortex shedding from the back of the building creates a temporary low pressure pulling the next vortex from the other side of the building into the wake. This low pressure, when time averaged, could be expected to channel contaminant up the back of the building. The higher you go, the faster and more violent the vortex shedding should be so the pressure deficit should increase to the top of the building. The Times Square simulation, already visualized in Figure 1, was instrumented with flow history tap locations up the back of the three buildings showing the effect. Pressure time histories were collected on these vertical lines but did not show any obvious lowering of pressure. This effect now appears to be driven by arch vortices lying behind the buildings – e.g., as in the well-studied problem of flow past a surface mounted cube (e.g., Martinuzzi and Tropea 1993). Further studies addressing issues of building geometry (i.e., aspect ratios) and angle of attack are needed to further characterize this important behavior.

Small-Scale Turbulent Fluctuations and Backscatter

The MILES subgrid-scale turbulence model built into FAST3D-CT has a desirable component of systematic backscatter actually incorporated implicitly (Fureby and Grinstein 2002). Stochastic backscatter arises when the phase-sensitive fluctuations at scales too small to be resolved numerically have a measurable effect on the larger, resolved scales. This component is usually not treated at all in LES models but plays a particularly important role in breaking symmetries and initiating turbulence.

Additional stochastic and systematic backscatter effects can be modeled by taking advantage of the flux-limiter information being computed by the FCT convection algorithms. This particular challenge is discussed by Patnaik, et al. (2005) in more depth than I will attempt here. An example of some preliminary computations for the stochastic backscatter a large city is given in that paper and the work for the 911-BIO program (Cybyk, et al. 1998) used this information in an algorithm designed to allow unresolved turbulence to migrate small particles to walls and other surfaces. These small-scale turbulence issues are on the forefront because physics, numerics and error analysis must all be rolled in together and then validated with experiments.

The Urban Geometry Challenge

Figure 10 shows the separate plumes at ground level from two short-duration contaminant releases at the locations marked by the yellow squares in downtown Chicago. The two problems were computed independently but simultaneously by FAST3D-CT using a single, chaotic flow field. The CFD mesh had 6-meter resolution in this simulation and was 360 cells x 360 cells x 55 cells high. The north-south separation between the sources is 350 meters, shown by the colored squares in the two panels. The wind is from the east (off the lake to the right) at 3 m/s. The conclusion is that the urban geometry has big effects that a model ought to capture. The differences between the two plumes must be attributed solely to the effects of the geometry. Channeling by the river appears to limit the spreading of the source in the panel on the left while recirculation and flow deflection behind buildings broadens the plume in the panel on the right.

Note that the right (northernmost) edge of the plume in the panel on the right, originating from the southern source location, reaches as far north as the plume edge from the northern source. This effect is seen in many urban situations. The geometry forms a natural boundary line that contaminant clouds can reach from a range of locations but have difficulty crossing. In this case the boundary seems to be formed by the drainage flow from the area just north of the Chicago River into and along the river. The source (yellow) for the plume in the panel on the right disperses quite a bit farther south, however, because this plume reaches the dense array of buildings just west of this source's location.

A reasonably faithful local geometry definition is required to ensure even the opportunity to capture the range of behavior illustrated in Figure 10. Unfortunately, detailed geometry descriptions of such large regions are seldom available. Regionally averaged building heights, densities and sizes are certainly not sufficient to capture the differences shown in Figure 10. Therefore, preparing a computer-ready digital representation of the terrain and building geometry is one of the important challenges and can be a showstopper. Simulating urban aerodynamics problems efficiently means that this geometry-definition process must be automated.

Overhead imagery, shown on the left in Figure 11 for a portion of an urban area, has been the traditional source of complex urban geometry. Processing stereo imagery, however, is a costly, lengthy process. Though techniques and software have improved considerably, more automation is possible and desirable. One can imagine a simpler

process to prepare a visual and computer-readable geometry representation of a large region where the height of trees, terrain, and buildings can be collapsed onto a single image, as seen in the figure. The panel on the right in Figure 11 shows one example of such a semi-quantitative geometry description. The representation uses different colors to indicate several ranges of building heights, different terrain heights, different land uses, and different tree heights. Such a representation is an easy way to characterize the geometry of a region and it permits, at least in principle, a better aerodynamic treatment than one based on a spatially varying roughness factor affecting the lowest layer of computational cells. Large steep and vertical surfaces are represented at least approximately to allow locally enhanced dispersion resulting from vortex shedding and building recirculation. The approach does not capture all the individual buildings and thus probably won't capture the striking differences seen in Figure 10 but does provide a way to represent large areas of buildings, for example a suburban subdivision, with a couple of parameters describing the size and density of the structures.

The horizontal resolution of the geometry representation desired is 5-meters to 10-meters, somewhat finer than the intended coarse-grained CFD simulations or very fine-grained meteorological models. Individual large buildings may appear but the representation is coarse enough that regions up to 50 km x 50 km can still be captured and viewed in a single file. Many CFD representations, however, require a rigorous geometry definition often characterized as "water-tight." This means that the edges of the polygons describing each body (building) must match perfectly with all the adjacent polygons. If the building were constructed according to the mathematical description, it would "hold water" with no leaks at any edges. Software would have to be developed to generate a watertight CFD grid from the simplified representation shown in Figure 11. This grid would embody only one realization of the geometry with the indicated characteristics. In our FAST3D-CT model a computational grid can be produced directly from a DEM representation such as shown in Figure 11 that is single-valued in height. As with the watertight grids, the representation is still just one particular realization. Every building would not be accurately sized or placed.

The advances of digital recording and laser technology now allow laser-imaging lidar over flights to collect raw data in the form of Digital Elevation Maps (DEM) at 1-meter to 3-meter resolution. Lidar data is already available for a number of regions, bases, and cities. The number and extent of these data sets is continually increasing and they are now often used, in conjunction with imagery, as the starting point to generate more complete surface/polygon representations of the geometry. Lidar datasets are generally adequate to specify detailed geometry even for the high-resolution CFD models but the digital elevation format requires costly preprocessing to provide shape files for the buildings, the polygon information required by most CFD models to generate a watertight grid. Furthermore, these preprocessing steps still do not properly capture the roof lines of many buildings or the locations and sizes of trees and other vegetation. Figure 12 shows such a lidar dataset for a small urban/suburban region.

Processing of lidar data generally begins with an edge detection algorithm based on height discontinuities to isolate building edges and walls from the ground. This edge detection also isolates dense blocks of trees that can appear to be irregular obstacles (buildings) with relatively rough edges. The distinction between buildings and trees must take two features into account, the extent of the region where the first and second lidar reflections ("returns") are appreciably different and the nature and quality of the fit to discontinuous height changes. In general, if the representation by a line discontinuity is very good, the line will be either a building pixel or the ground just outside the building. If the fit is poor and the variance of the difference between first and second reflections is large, the region may be judged to describe trees.

Identifying the various regions in the lidar image as being water, ground, trees and buildings does not complete the process because the height of the trees and the buildings has yet to be determined and this depends on the height of the ground under the buildings and trees. This can be solved as a constrained averaging problem, using the known ground and water heights as boundary conditions for a local averaging procedure for the undetermined regions in between. Once the ground height has been projected throughout the domain, the height of buildings and trees above the ground is determined by subtraction. The results of such an analysis are conceptually two arrays. The first is a pixel identification array suitable for plotting or modification into a geometry-registered map. The second is an array of heights associated with these "land use" identifications. The combined definition is sufficiently accurate, at least in principle, for any CFD solution resolving down to the street level.

The Emergency Assessment Problem

People who are responsible for making decisions in a public emergency involving airborne contaminants need fast, reliable assessments of the evolving situation, usually based on very limited data. They need help determining the

nature and location of the contamination sources and then reliable predictions of the path of smoke, particulates, toxic industrial chemicals or CBR agents. A quick response is required to save lives and to reduce the social and economic impacts of the incident. Methods and software to perform these assessments and the training necessary to take full advantage of them have become both military and civilian goals. However these requirements are expressed and satisfied, they must encompass the realistic features and time scales of a terrorist attack or an accidental contaminant release. Further, efforts to satisfy these requirements ought to leverage existing modeling and High Performance Computing (HPC) investments and provide an avenue to exploit breakthrough technologies as they become available.

Protecting a fixed site or region from airborne contaminants in real time is different from the careful, predictive simulation of a contaminant plume from a known set of initial conditions. The biggest difference is that very little may be known about the source, perhaps not even its location. Therefore analysis methods that are intended for emergency response should not require this information. The best we may have are reports of people becoming incapacitated, the existence of an unexplained traffic pile-up, or an isolated sensor detecting some contaminant at that sensor's location. It is crucial to be able to glean a useable situation assessment quickly from anecdotal reports, qualitative data, and any quantitative sensor data we may be lucky enough to have.

Waiting patiently for a computation to complete in the first minutes of an evolving disaster is both unreasonable and dangerous. Therefore software to be used for emergency assessment should be easy to use and invoke no computational delay. Delays of even a minute (not to mention delays of 5, 10 or 15 minutes) mean that the data on which the assessment is based are always out of date. The discussion of Figure 3 in the introduction interprets the human cost of this prediction delay graphically. We saw that 75% of the fatal doses are taken in the first 15 minutes in a typical urban scenario. Therefore new information and data demand immediate attention - including instantaneous computation of soon-to-be exposed regions and evaluations of the viable defense options. The system should also project optimal evacuation paths based on the estimated source location, winds, and predicted contamination regions.

In weather forecasting a long computational delay is dealt with by an expensive procedure called assimilation. New data at the current time are folded into ongoing predictions of tomorrow's weather that began hours in the past and have progressed well beyond the current time. In the very time-constrained defense against airborne contaminants, the only practical approach is to reduce the computational delay by a factor of 100 to 1000 from minutes to fractions of a second so the assimilation procedure used in weather forecasting is irrelevant.

Along with the ability for quick, accurate situation assessment, other emergency capabilities are required for operational use that are not strictly necessary for aerodynamic studies or detailed planning. Since each jurisdiction may wish to retain its current Geographical Information System (GIS) environment, it should be easy to import CBR assessments into existing systems and to broadcast the results in a graphical format that is easy to interpret and use. Finally, some additional advanced features are important. The software must accept weather information electronically where those services are provided, must accept and display remote sensor reports, and perhaps even integrate these various observations automatically. The next section shows one way to meet these requirements with research-grade aerodynamics computations.

5. Theory and Practice

The reasons to use Large Eddy Simulations to obtain high-fidelity urban aerodynamics solutions in lieu of other CFD models have been discussed. On the other hand, taking even a few minutes to solve a problem in a real world crisis is far too slow to be helpful. Faster than real time CFD computations have been possible for several years – though the quality of these solutions is suspect because of nonphysical assumptions including long time steps and coarse resolution. In theory, the continuing increases in computer power will allow more accurate CFD solutions to be performed fast enough for operational use in the future. However, even high-quality CFD solutions in a couple of minutes would not meet operational requirements. The inverse problem, finding an unknown source location based on measurements of the expanding cloud, is not easily solved using forward time integration. This is usually called the "backtrack" problem. The Navier-Stokes equations cannot be solved backward in time and approximate inverse procedures, such as computing adjoint solutions, are expensive and complicated. With enough computer power, a number of forward integrations could be performed to bracket an unknown source location iteratively but computer power has not advanced that quickly. Thus the problem of obtaining an operationally useful solution to the source backtrack problem has been a serious roadblock until recently (Boris 2001; Young, et al. 2004)

Therefore a two-pronged modeling approach has been adopted to date. Detailed models are run for planning and post-event forensics applications that are not time critical. Meanwhile, greatly simplified "common-use" models are used when speed is paramount or only small computers can be used. Constructing these common-use models is an art. Current versions have evolved from pollution prediction models and from battlefield models where the spatial scales and the time scales of interest are somewhat larger (longer) than is ideal for urban emergency applications. For civil defense and urban accidents, we are still faced with the usual dilemma: Accuracy or Speed? The remainder of this section describes one way to solve these operational problems almost instantly using high-fidelity CFD modeling.

A Practical Implementation Solving the Dilemma

An urban-oriented, emergency assessment system called CT-Analyst® has been developed for tracking airborne contaminants including Chemical, Biological, and Radiological (CBR) threats in deep urban geometries. Figure 13 shows a full screen image of the user display and Graphical User Interface (GUI) for CT-Analyst. The focus in this new urban capability is on the first few minutes to an hour after a CBR release and the first few miles from the source. Beyond these limits, the user has much more time to respond and changing weather data plays a progressively greater role. The example shown is for a ground-level release inside the red circle with the wind from 320^0 at 3 m/s. Figure 13 corresponds in detail to the FAST3D-CT simulation shown in Figure 4 for downtown Chicago. The geometry-dependent details on the right and left edges of the plume can be seen to correspond reasonably well on both figures. The regions of highest concentration and the general extent of the spreading also coincide closely. This figure also shows the ability of CT-Analyst to treat contamination reports and sensor readings, to estimate relative concentrations, and to display evacuation routes, the diagonal purple-magenta lines, optimized for minimum inhaled dose. This display takes less than 50 milliseconds to compute.

Splitting the operational problem into two stages defeats the accuracy-vs-speed dilemma. In the first stage, detailed CFD modeling of the domain generates solutions for a set of wind directions and source locations. This modeling is done off-line using FAST3D-CT. These engineering-grade computations use high-end parallel processing clusters. In this approach CFD experts in a laboratory environment execute the time-consuming portion of the prediction problem well in advance of an event. Because they are unencumbered by the time constraints of an actual crisis, they have the time and resources to review the output and resolve any numerical or resolution issues. They also have the option to adjust the resolution or the environmental parameters and re-run the model without losing time during a crisis.

In the second stage, the resulting ensemble of 3D CFD data is compressed using a patented compression technology called Dispersion Nomografs™. Nomografs compress gigabytes of urban airflow data by extracting the information important to first-responders and emergency managers. When deployed, the small Dispersion Nomograf database, a couple of megabytes per square kilometer, is loaded onto portable computers in the field so the system illustrated in Figure 13 is accessible to users in an operational setting. This paradigm expands the usefulness of CFD, making the results directly accessible to non-technical people and easily transferable among computers. When a contaminant release event occurs, the Dispersion Nomograf™ databases are displayed and manipulated using CT-Analyst. Since CT-Analyst doesn't integrate in time to generate new solutions, its response is immediate and its computational requirements are modest. Moreover, crisis management or law-enforcement personnel can master the CT-Analyst interface with a few hours of practice.

CT-Analyst has several unique new capabilities, gives greater accuracy and runs much faster than other current alternatives. This new, two-stage paradigm enables operational users to exploit the accuracy benefits of CFD-quality results without paying the costs of computational delay or extensive technical training. It also pays other important dividends such as the sensor fusion, sensor placement optimization, and the backtrack capabilities described below. Before discussing these enhancements I want to describe Dispersion Nomografs in a little more detail.

Dispersion Nomografs

A Dispersion Nomograf™ captures the distortions of plume shapes that result from the three-dimensionality of buildings, trees, and terrain - as determined by time-dependent CFD simulations that resolve urban geometries down to a few meters. The nomograf representation is actually somewhat more general. Nomografs can be constructed wherever the time-dependent, three-dimensional contaminant behavior is accurately known, i.e. from full fidelity simulations as in the current case, from other valid approximations, or even from ground truth data. Validation

studies of this new technology are being conducted and are described briefly in the section entitled ***Testing and Validation***. The plume "predictions" from CT-Analyst, based on a quantitative Figure of Merit, agree, within 80 to 90%, with the CFD simulations on which they are based.

Figure 5 and 6 show a number of snapshots of contaminant plume cross-sections in top view. Each of these plumes has a right edge and a left edge looking downwind from the location of the source, which is highlighted in the lower center portion of each of the panels. If another source were initiated at the right edge of one of the plumes, as indicated in the upper right panel of Figure 5, we would expect the new plume right edge (labeled R), indicated as a dotted white line, to follow the right edge contour of the original plume quite closely. The new left edge (labeled L) will follow it's own path across the original plume as shown. Dispersion nomografs provide a compact way to record this plume-edge information. The actual data format of the nomograf is chosen to allow very efficient algorithms for extracting plume footprint shapes corresponding to any source location or wind direction.

The nomograf data representation takes the form of two arrays of real values. An example of these two arrays is plotted in color in Figure 14, one array encoding the right edge information (looking downwind) and one for the left edges. One can view a dispersion nomograf as providing a coordinate transformation between the shape of the contamination footprint in a completely flat geometry and the shape found in the real world as determined by the detailed urban aerodynamics database. The plume envelopes within the contamination footprint and the concentration contours within these plume envelopes are implemented as simple, polynomial interpolations into the transformed footprint shapes. These polynomial representations are chosen for their simplicity, ease of computation, and fidelity to the trends extracted from the detailed simulations. The plumes move generally downwind and typically arrive on the ground from above because the contaminant generally moves faster above the ground.

Figure 15 shows the two sets of contours shown separately in Figure 14 overlaid on a simple rendition of the actual urban geometry being represented. Overlaying the left and right edge matrices, here displayed as contours, illustrates geometrically how the upwind and downwind sectors, originating at the intersection of one right edge contour and one left edge contour, can be defined for any particular location in the domain of the data structure.

A particular site, depicted by a square icon in the figure, represents a building. By selecting the left edge contour (upper lavender dashed line) and right edge contour (upper red dashed line) passing through this site, a danger zone for the site is defined in the upwind direction. Since the expansion of a contaminant cloud is contained within the cloud edge lines, any source outside of this shaded green danger zone, cannot reach the site in the current wind condition (shown by the downwind arrow). While the example of Figure 15 has a site in one particular place, users can place any number of sites at any locations within the region of the nomograf coverage.

The location of a contamination source is marked in Figure 15 with a star shaped icon. The source of contamination may be an accident, a leaking sprayer, or a broken container of hazardous chemicals that needs to be analyzed. By selecting the left edge contour (lower purple dashed line) and the right edge contour (lower red dashed line) passing through the source location, the source-processing algorithm identifies the contamination footprint shown in the lower right. A contamination footprint is the possible extent of contamination, given enough time to disperse throughout the region of interest. The actual contaminant plume envelope starts at the source near the upwind corner of the footprint and expands downwind away from the source as time increases. The plume envelope display is illustrated in the two panels of Figure 16.

Figure 16a, computed with realistic geometry, shows the predicted plume envelope at 6 minutes after contaminant release superimposed on the corresponding contamination footprint of a source (star icon). The figure also shows the upwind danger zone of a selected site (square icon). This result illustrates the various uses of the nomograf of Figure 14 and captures the contaminant flow using the full geometry of buildings, trees, and terrain. Figure 16b shows the corresponding plume envelope computed using a flat-earth nomograf with no building or tree geometry. Note the simple symmetry of the plume, footprint, and danger zone about the wind direction in this case and the qualitative differences of the results when the full geometry is considered.

Sensors are indicated in these figures with triangular icons. Blue sensor icons have not been triggered by contaminant at the time depicted in Figures 16a and 16b as they lie outside of the instantaneous plume envelope. Some sensors and sites will lie outside the footprint with flat-earth geometry but inside the footprint when full geometry is used. In other words, a prediction that does not account for complex geometry gives a false negative result for many locations, possibly resulting in unnecessarily high losses. An emergency response tool that does not account for complex geometry will be unreliable at best. Some conventional plume software overestimates the size of a plume envelope, particular near the source, to avoid these geometry-induced false negatives. Severe

overestimates can also have dangerous consequences. People may be inclined to stay in a dangerous location, thinking an uncontaminated region is too far away to reach.

The green shaded regions in Figures 16a and 16b estimate the upwind danger zone for the chosen site with and without capturing the effects of the geometry, respectively. The danger zone for a site is the set of all possible positions upwind where a source of contamination could reach the site. Computing the danger zone is an entirely new capability made possible by the nomograf representation and algorithms. The plume envelopes show the geographical region that the contaminant plume could have reached during its expansion up to the indicated time after release. The corresponding contamination footprints in Figures 16a and 16b respectively, represent the full extent of the growing contamination region after the plume envelope has spread to its maximum toxic extent.

Applications and Payoffs

New features are also available using nomografs – features that give first responders a chance to blunt a WMD attack rather than just understand the evolved catastrophe. Multiple sensor fusion is a consequence of the nomograf tabular form. Using three or four appropriate observations or sensor readings, CT-Analyst can backtrack to an unknown source location with zero computational delay. These capabilities depend on the pre-computation of the contaminant flow paths incorporated in the nomografs. They cannot be determined from forward integration of fluid-flow equations – no matter how fast.

Figure 17 shows a full screen view of the CT-Analyst tool in use for a large section of downtown Washington DC. The displays are overlaid on a map, familiar to anyone providing services within the city. Wind speed and direction can be set using the compass rose in the lower left-hand corner, by typing the numbers into the corresponding menu windows, or remotely through an Application Programming Interface (API). Basic display options and controls are located in the lower portion of the window. As shown in the figure, these displays have been activated, the Backtrack (dark and light blue regions in the lower, right portion of the window), the footprint (grey), and the Evacuation routes (purple diagonal lines). The scenario is the rapid release of a large amount of contaminant from a railroad tank car located at the star.

Figure 17 shows the Footprint display as modified by the Simulation time line, here set to 20 minutes after release. The Daytime density (concentration) display has been selected and accounts for the different color levels within the plume envelop display. The concentration model is three dimensional, conserves mass and is guaranteed positive. There is no magic in these simple mathematical approximations; they derive their accuracy from the fact that they are interpolations where the time integration has already been done carefully by the CFD. By way of contrast, even very complex, carefully contrived forward-integration schemes, with many effects and correction factors included, are still extrapolations in time in which small errors can accumulate quickly. Interpolation is always better.

Each location in a domain of interest has a downwind region as defined by Figure 15 (called the contamination footprint for sources) and an upwind region (the danger zone for sites). These two classes of regions are completely complementary, each being effectively the other's inverse. All assessments in CT-Analyst are "computed" by manipulating these two distinct regions for sensor report locations, for selected site locations, and for source locations. The dispersion nomograf representation makes these manipulations fast and requires only a minimum amount of data for each wind direction tabulated. The methodology can accept qualitative and anecdotal input and does not require knowledge of a source location or a source amount. In fact, backtrack to an unknown source, as shown in Figure 17, is accomplished graphically with zero latency using overlap operations on the danger zones of a number of "hot" and "cold" sensor reports.

One display button in CT-Analyst calls up the source backtrack capability, illustrated in Figure 17. The CT-Analyst backtrack can find an unknown source location based on sensor and observational data when the locally prevailing wind is known. The compound probability of the source location is computed by overlapping the upwind backtrack regions of all the active sensors.

To compute the danger zone, plume envelope, and backtracks displays knowing the actual concentration of the airborne agent is not needed. Indeed, until the total amount of the contaminant is known, plotting the actual concentration distribution isn't even possible. Therefore, CT-Analyst provides a relative concentration display until the mass of the agent from a specific source can be determined. Fortunately, this relative concentration and its time history provide enough information to minimize the inhaled dose of contaminant. The normalization used for the Chicago source in Figure 13 was chosen to correspond to the integrated mass of the source used in the FAST3D-CT

simulation of Figure 4. This normalization also accounts for the contaminant that leaves the grid through an analytic extension of the nomograf tables.

Visually comparing the CT-Analyst concentration plots in Figure 13 with the FAST3D-CT solution plotted in Figure 4 shows how well the nomografs capture the urban geometry-induced deviations from a smooth plume shape. Young, et al. (2004) describe and plot a quantitative Cumulative Figure of Merit designed to measure the congruence between a CT-Analyst solution and the underlying multi-realization database. For the case under discussion here the Cumulative Figure of Merit is about 80%.

The contamination footprints plotted by CT-Analyst are designed to provide plausible worst cases and are designed to "safe-side" the resulting situation assessments. The plume envelopes, which expand in time to fill the footprint, share this conservatism. This means that the edges of the plume envelope and the footprint are smoothed to maintain continuity so that the predicted contamination areas will be slightly larger than observed in the field. This is an interpretation designed for first responders. In practice this means that any particular realizations such as Figure 1 or 2, may only fill a part of the plume envelope depending on the structure of the wind gusts for that particular run. CT-Analyst attempts to indicate all regions that may be dangerously contaminated. This implementation is different from an ensemble average because the edges of the plume envelope are quite sharp and this is reflected in the concentration plots provided. This also is a feature of the individual high-resolution realizations during the first few minutes of any scenario.

Data analysis of different scenarios with multiple realizations using the Cumulative Figure of Merit support the following easy-to-remember interpretation. If you are outside the plume envelope predicted by CT-Analyst, you can be 95% sure of being in an uncontaminated region. If you are outside the contamination footprint, you are 98% certain of being in an uncontaminated region. These reasonable design goals give a simple way to express the uncertainty to a user or manager. Of course, if you have made a big error in the wind direction or the wind speed, some or all of the added fidelity will be wasted. Even in this circumstance, however, CT-Analyst seems advantageous because it so quickly allows evaluation of these real-world uncertainties.

CT-Analyst requires the current wind direction, speed, and atmospheric stability, e.g. time of day, whether it's cloudy or sunny, etc. This information could be input by the user or fed to the unit remotely. As the meteorological fidelity of the underlying 3D fluid dynamic simulations is improved, we can also improve the compressed CT-Analyst results by reducing the degree of conservatism and by providing analyses and displays that are more sensitive to meteorological factors. For planning future operations, CT-Analyst predictions are limited in accuracy by the quality of the wind forecasts provided for the operational area. The near zero-latency feature of CT-Analyst can be used to reduce the consequences of this uncertainty, however, by allowing rapid analysis of a wide range of probable conditions about those predicted in advance by the meteorological models.

As a planning tool, the CT-Analyst backtrack capability can be used to optimize sensor placement and thus reduce the cost of CBR defenses while increasing their effectiveness. Figure 5 in Young (2004) compares two sensor configurations for identical source and wind conditions. In one case eight sensors, six marking contaminant and two sensing clear air, were used to estimate the source location. The dark blue area shows the region of uncertainty for the source location, as shown also in Figure 17. The configuration and spacing of the sensors, the configuration and spacing of the buildings and the character of the wind fluctuations all enter the determination of this probable backtrack region. The comparison case has all conditions kept the same except two of the sensors have been removed and the four remaining sensors marking contaminant have been moved slightly. The result is an estimated backtrack region with half the area (half the uncertainty) and requiring 25% fewer sensors.

6. Testing and Validation

We test, verify, and validate CT models in four ways: compare the simulations with known solutions, compare with data taken in controlled laboratory experiments, compare with other well-tested models, and compare with data taken in field trials. The purpose of these activities is to build a basis of confidence in the model and to establish likely bounds on the errors that may occur in situations where the "correct" answers are unknown. Visualization serves a useful function in helping to elucidate the connections and relationships between the fluid dynamic phenomena, the geometry, and the contaminant distributions. There is certainly a subjective component to interpreting visualizations so their use can be subject to some abuse. However, things that look wrong are usually wrong, even if the converse is not always true.

The desired basis of confidence involves knowing with some degree of certainty that the model will work reliably in all regimes where it will be applied and that the errors have been quantified and are acceptably small. Obtaining this understanding is central to establishing a confidence in the model and the technology. Quantifying the errors is the difficult part because the variance in experimentally measured concentration values, particularly in urban-scale field trials, is generally unknown. CT field trials in cities are conducted to measure the airflow and simultaneously to determine the contaminant dispersion that occurs. The purpose is generally threefold: 1) to obtain ground truth about the flows and dispersion in particular regions of interest, 2) to study the fluid dynamic phenomena themselves in search of general truths to be applied elsewhere, and 3) to acquire data sets of sufficient depth and quality to calibrate and validate numerical models. Since modeling and simulation are the tools used to predict new situations, this third purpose sometimes appears to take precedence over the first two. Field trial experiments are large-scale, costly, and time consuming because a number of simultaneous activities must take place in different locations. Properly conducted field trials such as Urban 2000 (e.g., Allwine, et al. 2002) and Joint Urban 2003 take on the nature of an extended campaign rather than an experiment. Dozens of people come to town and may stay for weeks.

Determining and predicting pollution levels to determine air quality has been emphasized in the past. More recent field trials have studied airborne contaminant transport and dispersion related to accidents, industrial spills, and possible terrorist attacks. Measuring inert gas or particulate concentrations from known releases provides good data for planning civil defense because these are direct measures of the quantities of concern. Since the downwind contaminant plumes and quantitative concentration values are a complex integral over the fluctuating wind fields in a large volume, however, measurements that focus on tracer concentrations do not provide a good vehicle to validate dynamic flow models. Seldom are the local velocities in the urban canyons measured at enough stations and with enough time-resolution to validate a time-dependent LES model properly. Further, the number of separate trials is always limited and these are usually not repeated under statistically identical conditions to measure the naturally occurring variability.

Validation studies using field-trial data can determine whether the physics and fluid dynamics represented in a model are sufficient to describe the airflow and contaminant transport over a city. Field trials are always needed because only they can ensure that no surprises result from key physics left out of a model. Calibration is the process of using trusted data to adjust the free parameters in a model to give good agreement with that data. When the model being calibrated has key physics missing, however, it is still possible for calibration to give good agreement with the calibrating data set. The model will then give correct answers for the particular calibration data set but may still not predict other situations reliably. When dozens of variants of a numerical model are compared to data and the “best” one is selected, this also is a form of calibration and thus of only limited generality.

It is also possible to calibrate using incomplete or poorly characterized data sets. There is a tendency to characterize such calibration as validation once it is complete, but this is simply wrong. One key distinction between calibration and validation is tied to the ability to measure and control inflow and boundary flow characteristics and to repeat the experiments. Validation is a stronger statement than calibration because it entails a thorough characterization of the inflow and initial conditions for the experiments against which the computations are being compared. When it is not possible to say in detail exactly what problem a computational model should be solving, it is equally difficult to assess how well the model is actually doing. Adjusting various coefficients to get a better fit with incomplete data can cloud the issue rather than clarifying it.

Simpler models than detailed Large Eddy Simulations, such as RANS models, steady-state or time-implicit models, or Lagrangian puff models, have enough adjustable coefficients to calibrate using time-averaged data but their ultimate value for predicting new and dynamic situations is correspondingly suspect. Under the best of circumstances, validating urban airflow predictions with field-trial data can be comprehensive until the variance of the contaminant concentrations cannot be measured accurately in the field. Therefore other approaches are used to augment field trials for model validation in the years between field trial campaigns. Three of these alternate approaches are: 1) compare with known solutions, 2) compare urban flow simulations with controlled laboratory-scale wind tunnel experiments, and, 3) carry out model-to-model comparisons for simulations performed by different researchers with different models.

Validation of CT-Analyst and FAST3D-CT

The next few paragraphs discuss the validation of our FAST3D-CT and CT-Analyst models and provide a few examples of these four types of validation. Detailed analysis of these examples is beyond the intent of this paper.

Some additional details and references appear in the papers by Cybyk, et al. (1998), Boris, et al. (2001a, 2002, 2004, 2005), Patnaik, et al. (2003, 2005, 2006), and Young, et al. (2004). While the focus in this paper is on contaminant transport, the perspective adopted here is that no urban aerodynamics phenomena, including transport and dispersion, can be reliably predicted when the underlying flow physics is suspect.

FAST3D-CT validation has been underway for different applications for more than 30 years. Hundreds of publications record these efforts. The high-resolution Flux-Corrected Transport (FCT) convection algorithm underpinning the model (e.g. Boris 1971; Boris and Book 1973; Boris and Book 1976) was proved to be uniformly convergent (Ikeda and Nakagawa 1979). This means that the fluid dynamics in FAST3D-CT converges uniformly to the solution of the governing equations as grid resolution is improved, including real fluctuations and turbulence. Most other classes of T&D models do not have this property as they are actually applied. This means two things practically: 1) A detailed, time-accurate CFD model like FAST3D-CT has real predictive power and can be expected to be accurate for conditions significantly beyond those where it has been validated. 2) Physically-related validation studies in different problem contexts are relevant because the model being validated is predictive beyond the detailed regimes where it is validated.

FAST3D-CT and its predecessor models underwent testing and validation through the DoD HPCC CHSSI program (Lind et al. 1997) the BIO-911 ACTD program (Cybyk et al., 1998) and various Navy programs related to airflow over the superstructure of ships (Landsberg et al. 1995, 2000). The efforts have used standard benchmark tests, wind tunnel studies, and available full-scale data sets. Recognizing that a firm basis of confidence must include end-to-end studies of acute and even chronic contaminant release scenarios in urban areas we have recently been given access to the Joint Urban 2003 data set (Oklahoma City) and are beginning comprehensive studies of that data in conjunction with the University of Hamburg.

CT-Analyst is a look-up/interpolation technology that can be evaluated (verified) by how well it reports out the original data used to construct its Dispersion Nomograf™ databases. The CT-Analyst/nomograf technology was developed using FAST3D-CT but it could be extended to other simulation methods provided appropriate multi-source, time-accurate simulations are run and collected for nomograf preparation. In principle, with good data for an ensemble of experimental scenarios, a set of dispersion nomografs could be developed to capture the data set with a fidelity of 80 to possibly 90%. However, this model-to-model comparison (validation) procedure only indicates how closely the nomograf representation tracks an input data set. It cannot answer the question of how accurate the underlying CFD simulation methodology is. Data such as obtained from the Urban 2000 and Joint Urban 2003 field trials provide good *a posteriori* validation but blind, real time, in situ application of a model provides a better feedback about a nominally operational system the way it would actually be used. Greater confidence would result from computing the nomografs before the field trials and then to aid and interpret the experiments as they are happening and before they have been analyzed.

Generating the nomografs is a process of multi-variable calibration to the computed data for a particular region or city. This calibration is appropriate, however, because it is performed separately and in depth for each locale or city. Unlike calibrating a phenomenology using Salt Lake City data and then applying it in New York, separate customized nomografs are calibrated for each place to be treated by CT-Analyst. One would not use the nomografs for Washington DC to solve problems for Chicago. Therefore we can expect more accurate results than extrapolation technologies or calibrated models that are based on idealized obstacle configurations or data from different cities. Nevertheless, a strong basis of confidence requires that the nomograf encoding be compared to data sets for particular field trial locales as well to the CFD databases from which the nomografs are derived.

In the remainder of this section we will give a couple of short examples of the validation work underway. We must still acquire good data sets to validate some of the detailed sub-models in FAST3D-CT and CT-Analyst. We also need reliable data on local variability effects that influence the new sensor functions in CT-Analyst. As long as the models are being improved and the range of application is being extended, continuing validation work will be required.

Washington DC - Validation of Vortex Shedding in FAST3D-CT

Sometimes important validation comes in small chunks. Figure 18 shows a horizontal cross-section through a FAST3D-CT computation of the airflow over the Washington DC mall including the Washington Monument. The wind was steady at 3 m/s from the south and the cross-section shown in the figure was taken 20 meters above ground level. A very shallow boundary layer profile was imposed on the inflow so the velocity was essentially constant from near the base of the monument to its top. The variable visualized in the figure is the north-south (Y)

component of the flow velocity. Orange indicates faster flow than yellow and yellow is faster than green. Blue corresponds to regions of negative Y-velocity, such as found in building re-circulation zones. The wake of the Washington Monument is clearly visible as a sinusoidal region of momentum deficiency (green in the figure). The overlaid panel on the left side of Figure 18 shows an enlargement of the region north of the Washington monument from which the shedding wavelength and frequency can be determined. The monument is square, very tall, and quite isolated so this simulation provides a good, *in situ* test of the model's fidelity with respect to computing the vortex shedding necessary to compute dispersion. A lot of experimental data has been collected concerning the shedding of vortices from square and circular cylinders.

Figure 19 summarizes detailed computations of vortex shedding by FAST3D-CT and compares them with the experimentally measured values for both square and circular cylinders. The figure plots the non-dimensional Strouhal number for the shedding, versus the computational resolution of the structure. The Strouhal number is defined as the ratio of the characteristic wavelength in the wake to the size of the obstacle (thickness or diameter). A Strouhal number of 0.2 means the wavelength is five times the object's size. A resolution of 0.2 means that five computational cells resolve the object's diameter. The model was also run for ideal squares and circles at a wide range of resolutions in two dimensions. The horizontal bands in Figure 19 show the range of experimental values measured for squares (pink) and circles (green). The Washington monument data point, taken directly from the mall simulation shown in Figure 18, is indicated as a yellow ellipse in the figure. The ellipse axes indicate the estimated uncertainties in measuring the wavelength and resolution from Figure 18.

This is an important validation (verification) test because vortex shedding is the major contributor to contaminant de-trapping and dispersion in urban environments. Computing vortex shedding accurately, where it should be resolved on the building scale, is the best way to guarantee predictive fluid-dynamic treatment of the transport and dispersion and to maintain uniform convergence to the "correct" answer as resolution is improved (Boris 2002). Most other methods cannot solve this problem when run as they are used for contaminant transport scenarios. The dynamic vortex shedding effects are being replaced by a simple phenomenology such as a k-ε turbulence model or an externally specified diffusion coefficient. This unnecessary phenomenology introduces inaccuracies and means that improving the spatial resolution brings no additional accuracy. Averaging over important scales of motion artificially limits the accuracy. From Figure 19 we can see that the NRL CFD model would still shed vortices from the Washington Monument even at 10-meter or 20-meter resolution (one cell across the monument) but the frequency would be thirty to sixty percent low for such coarse resolutions. In the case illustrated here, FAST3D-CT comes within 3-5% of the expected shedding frequency when there are four cells across the obelisk.

Validation Using Wind-Tunnel Data on Urban Models

Comparisons with laboratory measurements of flow and contaminant dispersion through a simple urban model were made to evaluate and validate the ability of FAST3D-CT to model contaminant transport (Patnaik, et al. 2006). Brown et al. (2001) measured velocity distributions and tracer concentrations associated with the flow over an array of cubes in the USEPA wind tunnel facility under controlled conditions. The experiments were conducted in an open-return wind tunnel, with a working test section of length 18.3 m, width 3.7 m, and height approximately 2.1 m. The experiment approximated a neutrally stratified atmospheric boundary-layer flow over an array of cubical buildings. The array consisted of 7 x 11 cubes (0.15 x 0.15 x 0.15 m) with a one cube-height space between cubes. The velocity measurements, made with a pulsed-wire anemometer, consist of vertical profiles of the mean velocity and turbulence velocity variance in the three coordinate directions. These datasets are high quality and spatially dense but the data are not time-resolved. In addition to the velocity and turbulence data, the measured volume fraction for a C_2H_6 tracer released continuously on the centerline just behind the first cube was also reported. The laboratory profiles used as basic reference for the FAST3D-CT model benchmarking purposes were measured in the vertical symmetry plane of the building array.

Earlier studies using this data to test fluid dynamic modeling did not consider the effects of the CFD model on contaminant transport. Smith, Brown and DeCroix (2002) conducted such a study using the HIGRAD code, a CFD model that is second-order accurate in time and space with a Smagorinsky or a one-equation turbulent kinetic energy sub-grid closure. HIGRAD is normally used to predict atmospheric phenomena. Advection is done with the monotone MPDATA (NFV) scheme, making HIGRAD a MILES model. The HIGRAD simulations nicely reproduce the mean longitudinal velocity, including the recirculation patterns in the canyons behind the blocks. The turbulent kinetic energy is modeled well, except for some under prediction in the canyons. Another study by Lien and Yee (2004) modeled this USEPA wind-tunnel experiment on the cube array using a RANS STREAM code. The

Kato–Launder model was also used as an alternative to the standard Jones and Launder model. The agreement between the predicted mean velocity profiles and the experimental data is generally very good, with the greatest discrepancy occurring in the recirculation zone immediately downstream of the leeward face of the array. Their comparison indicates that a RANS model might be sufficient to predict the *mean flow* features.

Sometimes an open-air "wind tunnel" can be used. The MUST experiment, conducted at the U.S. Dugway Proving Ground in Utah, is one such series of field trials. An array of 10 by 12 shipping contains roughly 2.5 meters by 2.5 meters by 12 meters in size was laid out as an idealized scale model of a city. At this scale in the desert the natural boundary layer flow has about the same thickness relative to the containers as an urban boundary layer has relative to the buildings in a city and the shorter time scales allow several realizations of an experiment can be performed before the average environmental conditions will have changed much. The natural wind fluctuations, however, have relatively longer periods than in a city because the containers are about one fifth to one tenth building scale.

Time-dependent LES studies of this configuration were conducted by Camelli, et al. (2004, 2005) using a model called FEFLO-URBAN. They conducted a mesh-refinement study and wind-direction sensitivity study on this configuration. As generally expected in such turbulent flows, there was no indication of reaching a mesh-independent result even at the highest spatial resolution because more and more structure is resolved as the mesh is refined. To analyze the computations they defined a "region of interest," similar to the contamination footprint used by CT-Analyst, to capture the integrated effects of the highly intermittent flow. Their measures of congruence with the experimental results showed that 76% of the stations agreed with the experimental measurements within an order of magnitude. This number compares reasonably with the 90% number for FAST3D-CT comparison to Los Angeles field trial data presented in Table 1 below.

Houston TX - Comparing CT-Analyst with FAST3D-CT

Figure 20 compares contamination footprints computed by CT-Analyst with the evolving plumes simulated by FAST3D-CT for two release scenarios in downtown Houston. The "contaminant footprint" is defined as the overall potentially contaminated region and does not change in time, identifying the region that will be contaminated significantly, at least for some period during the event. The locally prevailing wind at 120 meters altitude is from the South at 3 m/s in this example with moderate wind fluctuations. The panels on the left, labeled a, b, and c, show a source located at the beginning of a relatively open area. A light blue star near the bottom of each panel indicates the location of the source. The panels on the right, labeled d, e, and f, show results for a source located near a number of large buildings. The time after release of the data in each of the three snapshots is indicated in each panel. Time increases from the top to bottom for each case. This is an example of verification of CT-Analyst by comparing two models.

In both scenarios the conservative nature of the contamination footprint provided by CT-Analyst relative to the CFD plume realization is shown. Once contaminated, a location stays contaminated. It overestimates the area by design to build in some assurance that the areas declared to be safe are actually safe. The footprint becomes more conservative the further downwind from the source location you look to account for the compounding effects of the uncertainty (variability) and the generally decreasing concentration of contaminant that is found far from the source. The footprint definition does not show the late-time reduction of the contaminant concentration, as seen, particularly in the panels on the left, from the FAST3D-CT computation. Therefore the dissipation of the FAST3D-CT plume in the relatively open area close to Source 1 in panels a, b and c is not seen in the footprint. The upper panels are shown 12 minutes after the release occurred and thus the region near the source on the left has already cleared out appreciably.

The CT-Analyst representation is doing a relatively good job of capturing the influence of the building geometry on the plume evolution. This is a validating result that can be communicated by visualization better than any other way. These two cases were the first two scenarios computed, visualized, and compared in this way with the corresponding FAST3D-CT footprints for Houston, so these probably are typical results. The nomografs were computed only once using the standard settings and with moderate wind fluctuations.

Comparisons between different code results for the same problem constitute validation when one of the codes, at least, is validated and known to be accurate. In the case described here and in Figure 20, FAST3D-CT is used to verify the compressed representation in CT-Analyst. Patnaik, et al. (2006) discuss another example of the comparison of two models, in that case FAST3D-CT and the Swedish FOI FOAM model. This form of bootstrap validation of model against model is generally not the preferred approach but takes on some merit as one or more of

the models involved in the comparison become recognized as reliable and accurate. This approach takes on great significance, however, when quantitative validation with field trial data is suspect because of the concentration variance problem. If two different LES models can be used to estimate the missing variance and these estimates agree, using this variance to compare simulated data sets with isolated field trial data sets becomes much more acceptable.

Los Angeles – Field Trial Comparisons with FAST3D-CT and CT-Analyst

Though modern field trials usually study long duration releases, the recent measurements made by Tracer ES&T, Inc in Los Angeles under U.S. Marine sponsorship featured short duration "acute" releases in a dense, deep urban area (Rappolt 2002). The depth and packing fraction of the Los Angeles buildings in the downtown area is typical of large cities and the short duration releases are closer to expected terrorist scenarios than the chronic (pollution-like) scenarios featured in earlier large-scale exercises. Eight different realizations of FAST3D-CT simulations for one of the LA field experiments are shown in Figure 5. These were computed for moderate wind fluctuations impressed at the upwind boundary. Figure 6 shows the corresponding eight realizations computed with steady-wind boundary conditions (zero fluctuations impressed). Figure 21 shows the 3rd and 5th realizations of Figure 5 enlarged so the experimental values plotted within the colored squares are more easily seen. The release point in each panel is the white circle and the dark blue squares indicated experimental values below a 20 PPT threshold. The FAST3D-CT results for the two realizations are seen to be quite different and yet each seems to agree reasonably well with the experimental sampler values shown in the squares. The next few paragraphs consider the quantitative aspects of this apparent quandary. How different can two solutions be and still agree, within tolerance, to a given data set? Alternately, how much stock should be placed in the particular values of one particular data set?

The Tracer short duration "acute" experiments consisted of SF6 being released continuously for five minutes in each of twelve widely-separated trials. 50 synchronized samplers took 12 gas collection samples, each of 2.5 minutes duration, for a total experimental trial time of 30 minutes after release. The region instrumented was about one kilometer square as shown by the square sampler locations in the panels of Figures 5, 6, and 21. These scenarios focus on the first thirty minutes when prompt action will be most effective in reducing loss of life. FAST3D-CT was run for the same conditions as the trial (as closely as could be determined) with wind fluctuations imposed at a moderate level and the sun set to mid morning for the particular trial (#8) chosen. Eight independent realizations of the release were computed in ambient wind conditions taken from the same synthetic distribution of fluctuations. These correspond to releases five minutes apart in a continuously computed flow field. Five minutes was measured to be an adequate de-correlation time.

As described briefly earlier, cross-sections of the contaminant concentration for all eight numerical realizations are shown in Figure 5 for the tenth sampling interval, 20.0 to 22.5 minutes. This interval begins 22.5 minutes after the SF6 began being released at –2.5 minutes. The differences from one realization to another are substantial. For example, look at the first, third, fifth and seventh panels in Figure 5. It is reasonable to conclude that there can be no single correct experimental data set to be compared with simulations or models – even for perfect field trials conducted under environmental conditions identical to those of the simulations. Measuring multiple realizations of the city-scale experiment is usually not possible before conditions change appreciably so there is no quantitative yardstick for comparing measured and computed concentrations.

The ensemble of multiple simulated realizations, however, can be used to define the missing variance and thus it provides the necessary "yardstick" to make quantitative comparisons when the computational conditions, including the wind fluctuations as input, are well matched to the experimental conditions. When wind gusts are turned off all together, as in Figure 6, only a small reduction in the naturally occurring concentration variability is observed because building vortex shedding still provides appreciable turbulence and the neutral boundary layer in the simulations is unstable from solar heating.

Several of the accepted techniques for comparing a simulation data set with an experimental realization, as described by Chang and Hanna (2004), were applied to these data. Congruency counts are one of the methods recommended. For the eight computational realizations of the baseline CFD simulation (wind from 170^0 at 3 m/s with moderate wind fluctuations) the percentage of simulation data points within 20% of the experimental values, the percentage within a factor of 2, within a factor of 5, and within a factor of 10 were computed. The following table summarizes these results for the ensemble of eight realizations in the baseline simulation.

Table 1. Congruency Counts for Field Trial #8 and Simulation Data Set Comparisons

Approximate Wind Condition: 170^0 at 3 m/s (fluctuating 50%)	+/- 20%	Within a Factor 2	Within a Factor 5	Within a Factor 10	Counts
Field Trial #8 vs Simulation Set	14.5%	39.5%	73.8%	88.8%	1147
Realization 3 vs Simulation Set	19.5%	53.6%	84.4%	90.0%	1496
Realization 5 vs Simulation Set	15.7%	46.9%	81.9%	93.0%	1161

Only about 160 of the possible 600 experimental values were above the threshold of 20 parts per trillion for Trial #8. The total number of counts in Table 1 is about eight times this number because there were eight CFD realizations. It is seen that almost 90% of all observations were within a factor of ten and almost 75% were within a factor of five, but barely 40% were within a factor of two. Only about 15% of the number pairs were within 20% of each other when the experimental values from Trial #8 were compared to the simulations. Is this good or bad agreement? Judging from studies on Salt Lake City sampler data, this is not very good agreement. Those comparisons, however, were for gas samples averaged over half an hour, not 2.5 minutes, and for distances of several kilometers from the source, not for distances less than a kilometer.

To assess the importance of the large concentration variability revealed in Figures 5 and 6, this congruency test diagnostic was used to compare a numerical realization, as if it were experimental data, with the ensemble comprised of the remaining seven realizations. Thus no particular realization was compared with itself. The two concentration distributions, Realization 3 and Realization 5, shown in Figure 21 for the 10^{th} sampler interval were used for this congruency test. We know that each of these realizations comes from the same distribution as the ensemble so the lack of perfect agreement in the bottom two lines of Table 1 is a result of the natural variability from one realization to another. This congruency pattern is not the result of errors in the solutions or systematic difference between the experimental and simulated wind fluctuation distributions. Indeed, the results given in Table 1 are quite typical of simulated ensembles. Here, less than 20% of the number pairs are within 20%, about 50% are within a factor of two, less than 85% within a factor of five, and again about 90% within a factor of ten. The experiment-to-simulation congruency pattern is almost the same as for the simulation-to-simulation comparisons. Indeed these may be statistically indistinguishable. Further, there is little reason to expect that closer agreement is possible between simulation and experiment, given the differences from one realization to another.

To obtain a quantitative algorithm to make these comparisons, we studied the distributions of the simulated SF6 concentration values in the ensemble of realizations in order to compute the concentration variance numerically. Figure 22 shows the distributions of simulated concentration values near Sampler 25 for each of the twelve sampling intervals in Trial #8 beginning with the SF6 release (T01 in the upper left) and continuing for half an hour (T12 in the lower right). "rlz = 9" indicates the composite distribution of concentration values from all eight realizations. Sampler 25 was chosen for this figure because it was quite close to the source and so measured concentrations above threshold for all 12 sampling intervals. The horizontal concentration scale is logarithmic, as used in other validation studies. Numerical concentration values are collected at the experimental sampling sites and at nearby points in the simulation to build up the relatively continuous distributions shown. With the exception of the first couple of sampling intervals, these concentration distribution functions are seen to be reasonably normal in shape so computing the mean and standard deviation makes sense. This standard deviation shown by the horizontal bar approximates the concentration variance needed to compare simulated and experimental values quantitatively.

This computed variance is derived from the eight different realizations and is illustrated by the horizontal black bars in each of the 12 distributions plotted in Figure 22. It gives a point-by-point yardstick to compare simulations and experiments, a metric lacking in single experiment field trials. The center of the horizontal black bar in each panel is at the mean value of the distribution and the bar extends one standard deviation on either side of the mean. We define $X^2(x,t) = (N_E(x,t) - N_S(x,t))^2/\sigma(x,t)$, the normalized difference of the simulated and the measured values. Since X(x,t) is a normally distributed random variable with unit variance, $X^2 = \Sigma_v X^2(x,t)$ is said to obey the chi-square distribution with n degrees of freedom where n is the total number of independent x and t values at which meaningful experiment-simulation data pairs can be defined.

The Chi-square value for the set of twelve measurements indicated in Figure 22 is ~13 and the corresponding probability of agreement is 45%. This particular statistical comparison shows that the experimental data taken by sampler 25 has a 50-50 chance of having been drawn statistically from the solution space defined by the ensemble of simulations because the chi-square value is about 13 for 12 degrees of freedom (Abramowitz and Stegun, 1964). This agreement is not very good because the sampler was very close to the source. Most of the other samplers that

were above threshold for a few of the sampling intervals showed much better agreement using the chi-square test. When all samplers and time intervals were taken into account, there were 159 degrees of freedom in the experiment and the baseline simulation ensemble showed a 98% Agreement Probability with the field trial #8. This means that this level of agreement could occur by chance less than one time in fifty. This analysis approach appears to be somewhat more capable than other analyses of saying something quantitative about the validation of the detailed simulations with a single data set.

The results to date show that agreement between the experiments and a set of distinct FAST3D-CT realizations are as close statistically as can be expected, given the strong variability of the flow field environment. Chi-square tests of agreement range from 98% to 99.9%, as discussed further with respect to Table 2 below. Better CFD computations may be possible than those presented here, but it is hard to see how we could determine this, given the existence of only one field trial data set for each condition and the large natural variability. Use of the community-accepted comparison measures show a comparable quality of agreement but these approaches are less sensitive in distinguishing small parameter variations such as wind speed and direction (see below).

Because the CFD results and the CT-Analyst nomograf representation distilled from the CFD runs can both be compared with the experimental data, these comparisons qualify as an end-to-end validation effort. *In situ* end-to-end validation of the entire procedure using the field trial data is possible using the CT-Analyst concentration predictions. Figures 23a and 23b show visualizations of each of the twelve time interval data sets for Trial #8, used in the FAST3D-CT comparisons above, superimposed on the corresponding CT-Analyst plume concentration plots. The first six sampling intervals are shown in Figure 23a and the last six in Figure 23b. The colored squares are the actual field measurements overlaid on the CT-Analyst concentration for the corresponding scenario at the same time. Although the color scales are not identical, again we can see that the agreement is good.

The actual threshold value used for the experimental measurements lies somewhere in the middle of the concentration range indicated by the dark green regions near the edge of the plume envelope. Another analysis of this comparison can be performed by counting the number of sampling locations where CT-Analyst over-predicts contamination, locations where CT-Analyst and the data agree that there is contamination above threshold, and locations where the data indicate contamination but CT-Analyst does to predict any at the specified time.

A number of the sampler locations are crossed out (red X) where the experimentalists felt the data was invalid. Some of the short duration, out-lying points perhaps should have been declared invalid. Taking the data as they are for the twelve time slices, we can apply a form of qualitative analysis designed for such comparisons (Warner, et al. 2001). About 80 data points are over-predicted (conservative plume envelope and concentration contours), about 145 simulated data points are in agreement with the trial, and about 15 data points are under predicted. Of the 245 "warnings" that would be issued in this data population about 70% are true and 30% false positives. Of the "exposed" population (either predicted and measured or measured) consisting of 160 data points, 91% are properly evaluated and 9% are false negatives (i.e. exposed but not warned). We can use these values to compute the guaranteed assurance percentage described earlier (nominally set at 95%), 15 counts are outside the 245 warnings issued and correspond to the exposures that would be penalized in the Figure-of-Merit computations described earlier. This corresponds to a 94% assurance that if you heed the CT-Analyst predictions you will be safe.

There are no surprises here and no manifest problems. Returning to the visualizations and its interpretation for Figures 23a and 23b, it is clear that most of the false positives (over predictions) occur on the north and east edges of the advancing plume. This further corroborates the observation, discussed in the next section, suggesting that using a wind from about 165^0 at about 2.5 meters per second might be a better match to this particular data realization. It could probably reduce the false positives somewhat while not greatly increasing the false negatives. This remains to be seen. In fact, it is unclear whether the locally prevailing wind direction and wind speed can be measured reliably to within 5 degrees and half a meter per second respectively in a real emergency situation.

Parametric Variations

Using model runs to determine or refine unknown parameter values is a well-accepted procedure, e.g. (Wirsching, Paez & Ortiz 1995; Bendat & Piersol 2000). Therefore, we are entitled to use the chi-square data-set comparison mechanism to test a number of physical parameters and sensitivities quantitatively. Table 2 below summarizes seven of these chi-square tests of agreement between the Tracer ES&T, Inc. data set for Los Angeles Trial #8 and CFD MILES ensembles run with physically different parameters. In each case the experimental data set is the same but the eight realizations change as the system parameters are changed. Case 1 is the baseline case.

Table 2. Studies of Sensitivity to Wind Parameters	
Case# (all 159 DOF)	**Agreement Probability**
1. 170 deg @ 3 m/s (100 m radius)	**.982 (baseline)**
2. Low Wind Fluctuations	**.963**
3. Steady winds (buildings shed)	**.935**
4. 160 degrees @ 3 m/s	**.997**
5. 180 degrees @ 3 m/s	**.371**
6. 170 degrees @ 2 m/s	**.999**
7. 170 degrees @ 4 m/s	**.023**

Case 2 was run with lower wind fluctuations than the baseline case (Case 1) and Case 3 had no superimposed wind fluctuations, as discussed and illustrated with Figure 6 in Section 2 and discussed above. Even though the inflow wind is steady in Case 3, the city rapidly generates fluctuations from building vortex shedding and uneven heating. In both these cases the chi-square probability of agreement is high, 96% and 93% respectively, but not as high as for the baseline case. These three cases suggest that the baseline case is a better match to the data – but not by much. Using this procedure to try to pin down the wind fluctuation strengths and scale lengths, however, may not be a good idea. The amplitude of the wind fluctuations ought to come from a measurement source. It cannot be determined from this procedure self-consistently because the variance used to compare simulation and experiment increases as the wind gust strengths increase. The agreement between the simulated ensemble and the experimental data set is bound to look better as the yardstick itself gets longer, i.e. the expected variance is increased.

Cases 4 and 5 are simulation ensembles where the wind direction was varied from the nominal baseline run at 170^0 at 3 m/s in Case 1. From the results, having the wind from 180^0 at 3 m/s gives unacceptably poor agreement. However, 160^0 at 3 m/s actually appears to be in better agreement with the experiment than the baseline case. For comparison, when we used the congruency test as given above for these different simulations, it could not be determined with any confidence that something other than the baseline case might actually be better. The last two cases, 6 and 7, are different analyses of the baseline case wind direction with different speeds. The agreement gets worse when the average speed is increased to 4 m/s and gets better when the speed is reduced to 2 m/s. These crude sensitivity studies are just scratching the surface here. The data processing is relatively straightforward so runs at intermediate parameter values are clearly indicated to establish how broad the parameter region of optimal agreement actually is.

Meanwhile, an important conclusion should not be overlooked here; better computations may be possible but it will be very difficult to determine that they are better because of the very large concentration variability to be expected in any field trial data for acute releases. Problems that will always attend such comparisons include: what level of contamination should be chosen as the threshold, exactly how to define the contamination footprint and plume envelope (given that CT-Analyst can compute a relative concentration), how to display variability, and what levels of assurance are appropriate for general use. For this latter question: we are attempting to implement 95% assurance that being outside the plume envelope means you are currently uncontaminated and 98% assurance that being outside the footprint means you won't be contaminated. Of course, these estimates assume the "correct" wind direction and speed have been specified to the CT-Analyst system.

To summarize our comparisons with the Los Angeles field trial data:

1) Appropriate, multi-realization FAST3D-CT simulations with 6-meter resolution seem to be virtually indistinguishable from the Tracer Los Angeles field trial data using the chi-square.

2) The natural difference between realizations can be quite large, due to building vortex shedding, even when inflow wind gusts are turned off.

3) Multi-realization CFD simulations provide a way to approximate the missing variability data and thus enable quantitative comparison of data sets using the chi-square test of agreement.

4) The chi-square probability also gives a sensitive way to check congruence, to approximate unknown parameters (Bendat & Piersol, 2000, Wirsching, Paez & Ortiz, 1995), and to validate time-dependent CFD models.

5) Direct comparison of the CT-Analyst predictions with the experimental data set shows some over prediction, as expected and desired, and also corroborates the design goal of 95% assurance that being outside the plume envelope at any time means you are actually in an uncontaminated region.

This discussion of the Los Angeles trials summarizes preliminary results only. Considerably more qualitative and quantitative analysis will be undertaken on the other trials available. We are also following up with comparable analyses and comparisons in another major urban field trial situations (complex topography) where wind tunnel data as well as field trial data are available. The scientific validation community for CFD, resident in the various professional engineering societies, universally recognizes the need for multiple experimental data sets (realizations) to establish quantitative measures of the variance in experimental data. Field trials cannot provide this reliably but wind tunnel studies on the urban geometry and multi-realization simulated data sets can.

7. Conclusions

Any approach to situation assessment for airborne contaminant emergencies is based on an interpretive transport model, whether or not this is explicitly stated. The model may be only a set of qualitative notions in the head of an operator or commander that contaminants will generally blow downwind at some rate and spread a lot. Usually this interpretive model is at least in part computational. In a crisis people want to know they are making good decisions based on a good understanding of the evolving situation. In the past, unfortunately, more accuracy has always meant more computing and more computing means more delay. Waiting even one or two-minutes for each approximate scenario computation can be too long for timely situation assessment. State-of-the-art engineering-quality 3D predictions, that one might be inclined to believe are accurate, still take hours or days.

This paper considered urban aerodynamics challenges in two classes: problems where there is enough time to prepare a high-resolution solution using significant computer resources and problems where good answers are desired but where instantaneous answers are required. A new paradigm was presented for bridging these challenges by bringing CFD-quality solutions to time critical emergency decisions. In addition, an operational example of this paradigm was presented, NRL's CT-Analyst system. Table 3 summarizes the several advantages of using this new paradigm, which requires substantial computing time but performs this computing well before any time-critical applications of the tool. The waiting time for the crisis manager to obtain useful results is greatly reduced, his training time is much shorter, the system data storage requirements are less, and the number of lives that can be saved is greatly increased. As can be seen, CT-Analyst is about a thousand times faster than alternatives and can be learned in one percent of the time.

Table 3. Comparing Airborne Contaminant Plume Solution Methods

	Approximate Time to Run (seconds)	***Approximate Time to Learn to Use (days)***	***Approx. Data Requirements (MB)***	***% Lives Saved***
CFD Models	**50,000 (1 hr/scenario)**	**1,000 (~Ph.D.)**	**180,000**	**0% Too Slow**
Puff/Plume Models	**50**	**10 (1-2 weeks)**	**180**	**50% (~10 min)**
CT-Analyst Using Nomografs	**0.05**	**0.1 (~2 hours)**	**18**	**85% (3-5 min)**

The current common-use models are much faster than CFD but still can take a minute or more to run – once all the data have been entered for the scenario via pull-down menus and dialogue boxes. CT-Analyst computes and displays each scenario another factor of one thousand faster. This speed does not come at the expense of accuracy and is great enough that multivariable sensor optimizations can be performed directly in hours where the corresponding computations with other models would take months or years (Obenschain, et al. 2004).

Concerning Operational Models for Defense Against Airborne Contaminants

The technical questions that first responders and agency heads alike want to have answered are: What speed is really necessary? What capabilities are required? How much accuracy is enough? Underlying these questions is an even deeper one: What good can a first responder expect to do with even optimal emergency assessment tools?

Figure 3, which is typical but by no means comprehensive, answers this last question in terms of lives that can potentially be saved. Put in plain terms, five out of six people who would die without any warning of the inhalation danger can potentially be saved by a timely warning that is accompanied by the correct direction to walk. This warning must identify the plume centerline and councils evacuation away from the centerline across the wind. This 15 percent fatality rate should be compared with the 75 percent fatality rate that results with evacuation or sheltering-in-place when the warning takes 15 minutes to be issued. Every minute of delay in issuing the warning translates into an added 5 percent fatality rate.

These results directly affect how the first three questions should be answered. To assess an emergency situation in the first few minutes after a contaminant release requires a number of model evaluations. Performing these model runs in a few seconds each would be adequate in a crisis as long as the total set-up and run time was only a couple of minutes. If the initial knowledge of the conditions of the event were perfectly known, three transport and dispersion runs might be adequate, one for the nominal scenario and two more bracketing the estimated wind direction by plus and minus 10 degrees to estimate safety factors against unforeseen wind changes.

In situations such as accidents or terrorist attacks, however, the parameters of the release are not likely to be known and will have to be guessed or approximated. This requires a number of additional runs. Backtracking from a few observations, for example, requires the computational equivalent of at least one run for each sensor report or observation in the backtrack computation. Iterating forward plume computations and then choosing the source location that best fits the available data can also approximate the CT-Analyst backtrack procedure. This iterative procedure will take at least a few runs to execute and requires some quantitative algorithm to select the best solution. If each of the individual runs takes even a minute, the overall procedure will take up to ten minutes. This is too long.

The speed of a CT model can be even more important when it is being used for planning purposes. Because so many more runs should be performed, the total computer running time can be more of a limitation for planning than for the analysis of emergencies. In system design and optimization, tens of thousands or even millions of runs are required. Studies conducted to look for weak points in the defense of a facility demand extensive parameter surveys. Typically hundreds of individual plume evaluations are needed for each wind angle and each different type of source. Jointly optimizing the placement of the sensors in a detection network is yet more demanding. In all but the most trivial cases the only way to guarantee an optimal configuration is by exhaustive search (Obenschain, et al. 2004). Millions of sensor coverage evaluations are needed to evaluate and compare fitness functions for entire sensor network configurations. To accomplish the optimization in a few hours with a single processor requires individual evaluations in 10 to 50 milliseconds. If each plume or coverage evaluation takes a minute instead, a single processor would have to run for about two years.

Therefore crisis management, combined with planning applications, suggests what particular capabilities should be elevated to requirements. For crisis management a backtracking capability saves a lot of time and this time can result in saved lives. Even when multiple forward runs can be iterated against observations and sensor data, having a composite backtrack function to perform the information fusion will be less cumbersome, easier to use, and less subject to human error in times of stress. Completely automating the procedure is not necessary and possibly not desirable. Even when complete automation becomes possible, there is still the issue of trusting the individual items of input data. Are the sensors sensitive enough? Did the person pass out from something other than an airborne contaminant? What about false alarms? A partially manual system with an operator in control allows the distinct observations and reports to be individually included in the computation or excluded, giving a means to assure the crisis managers that the conclusions are robust against misinterpretations of the data.

It also makes sense to structure the software to compute and display what a crisis manager really wants to know. Evacuation routes are simple to determine, in principle, when the geometric effects are negligible but the geometric distortion of a contaminant plume by buildings and terrain can have a large effect on the optimum direction to go from each location in the domain. Such capabilities should be extended to compare evacuation and sheltering options in terms of relative doses of contaminant that would be received. This latter capability implies that concentration curves can be evaluated at the locations of all the relevant buildings and the ordinary differential equations that describe how the contaminant infiltrates the building can be integrated essentially instantly.

To conduct planning surveys efficiently suggests developing a capability to execute many scenario evaluations and to analyze the results automatically. A scripting system or automatic scenario generator is required. A general way to accomplish this has been designed into the CT-Analyst Application Programming Interface (API). The accuracy of CT-Analyst, coupled with the capability to perform thousands of scenarios in an hour, makes the system ideal for site defense, personnel protection, battle management, and operations planning as well as detailed forensics. The speed and accuracy make the system well suited to virtual reality training.

As a planning tool, the CT-Analyst backtrack capability can also be used to optimize sensor placement and thus reduce the overall cost of WMD defenses while increasing their effectiveness. The objective is to save money while increasing overall system effectiveness by reducing the number of sensors needed and increasing the utility of the already existing ones. If successful, this is certainly an important potential payoff. Figure 5 in Young, et al. (2004) compares two sensor configurations for identical source and wind conditions. One configuration was taken from a standard scenario used for demonstrating CT-Analyst. In this case eight sensors, six hot (marking contaminant) and two cold (sensing clear air), were used to estimate the source location. In the comparison case all conditions were kept the same except two of the sensors were removed and the four remaining sensors detecting contaminant were moved slightly. The result is an estimated backtrack region with half the area (half the uncertainty) that requires 25% fewer sensors. If several tens of millions of dollars were being spent on the equipment for a city, such computational analyses could save five or ten million dollars.

Viewed as an optimization problem, placing a number of sensors in a coordinated network is not soluble in closed form for any but the most trivial cases. The only guaranteed optimization procedure is exhaustive enumeration of all possible sensor placement patterns. The configuration and spacing of the sensors, the configuration and spacing of the buildings, and the character of the wind fluctuations all enter the determination of the optimal sensor placement pattern in the real world. This important procedure is considered in considerably greater detail in Obenschain, et al. (2004) where a parallel genetic algorithm was employed, enabled by a new sensor coverage function in CT-Analyst, to jointly optimize the placement of 40 point detection sensors in the downtown area of a real city. This computation took four hours on eight processors to perform, evaluating about one million configurations, each with 40 sensors. The equivalent computation using existing plume models would have taken two years or more. This study showed that about a factor two increase in the sensor utilization efficiency over simpler approaches was achievable. In addition the effects of the city geometry are directly factored into the optimization procedure without additional cost.

Concerning Detailed Urban Aerodynamics Models

The accuracy required of operational models is certainly less than required of the scientific models and one should not mix up these uses. Several questions spring to mind. Is it worthwhile working this hard? FAST3D-CT is fast and easy to use but still requires work to use properly and even more work to improve. Should the CFD community be working harder – considering that it is very difficult to measure just how well you are doing? If working harder seems worthwhile – what should we be working on? These are important questions because the natural variability and uncertainty are so large that many people have been tempted to argue that we should not worry too much about the detailed accuracy of the prediction models.

My viewpoint is that we should be working hard enough to guarantee that the variability and uncertainty in the wind parameters are the main errors – and be prepared to show this. These seem to be the largest uncontrolled parameters in the solutions and they set the levels below which other errors may be undetectable. There are a number of processes and factors that go into making a detailed urban aerodynamics model. If even a number of these processes are allowed to approach the inaccuracy of the variable wind information, the composite errors would often be much larger. Further, ensuring that only the wind variability (uncertainty) is a major source of error means that sensitivity studies performed with the model will be informative and reliable even though the specific cases simulated may only be representative.

Necessary detailed model improvements include more than just testing and validating what has already been done. We have just scratched the surface of buoyancy and heat transfer models. The inflow temperature profiles contain the effects of the meteorological heating and heat transfer upwind but this heat transfer must be maintained consistently across the entire computational domain.

The high-fidelity external airflow should also be connected to building infiltration models and to subway system level models. It is important to improve the land use representations and the uses that the numerical models can make of this information. By the same token, physical coefficients for many of the effects that are sometimes important need to be pinned down. Turbuophoresis is the transport of particles and gases to surfaces in situations where the average velocity is zero and thus the macroscopic convention is insignificant. It deposits particles on ceilings as well as floors. This physical effect seems an ideal application of the stochastic backscatter analysis made possible by the form of the MILES subgrid turbulence model.

Sensitivity studies need to be performed for a number of physical effects in an academically supportable manner. New physical effects need to be put in the detailed models and the changes in the solutions they cause should be assessed. For those where the effects are important in significant situations, the models need to be inspected and upgraded to ensure the errors are smaller than the effects of wind uncertainty and wind variability. Then the tests need to be performed again.

I would like to close with an important conclusion about the whole process. Better computations may be possible but it will be very difficult to determine that they are better because of the large variability to be expected in any field trial data for acute releases. Problems that will always attend such comparisons include: what level of contamination should be chosen as the threshold, exactly how to define the contamination footprint and plume envelope, how to display variability, and what levels of assurance are appropriate for general use? For this latter question the NRL team is attempting to implement 95% assurance that being outside the plume envelope means you are currently uncontaminated and 98% assurance that being outside the footprint means you won't be contaminated.

Acknowledgments

The author wishes to thank Bob Doyle, Adam Moses, Keith Obenschain, Gopal Patnaik, Julie Pullen, Lisa Williams, and Theodore Young, Jr., the NRL contaminant transport team who helped to make this technology possible. Further thanks are due Carey Cox, Bo Cybyk, Jack Fulton, John Iselin, Sandy Landsberg, Charles Lind, Warren Schultz, and Rob Scott for their technical discussions and contributions to various components of this effort. Aspects of the work discussed here were supported by ONR through NRL, the DoD High Performance Computing Modernization Office, NSWC Crane, DARPA, DTRA, TSWG, and the MDA Post Engagement Ground Effects Model (PEGEM) effort.

References

Abramowitz, M. and A.I. Stegun, 1964, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, NBS Applied Mathematics Series 55, U.S. GPO, Washington DC, pp. 940-943.

Aliabadi, S., A. Johnson, J. Abedi, S. Tu and A. Tate, 2004, "High Performance Simulation of Contaminant Dispersion on the Cray X1: Verification and Implementation", *Journal of Aerospace Engineering Computing, Information, and Communication*, **1(8)**, 1542-9423, pp. 341-361.

Allwine, K.J., J.H. Shinn, G.E. Streit, K.L. Clawson, and M. Brown, 2002, "Overview of Urban 2000: A multiscale Field Study of Dispersion through an Urban Environment," *Bull. Am. Met. Soc.,* April 2002, pp. 521-536.

Baffour, R., S. Aliabadi and A. Ji, "GIS-Based Complex Geometric Modeling for High Performance Computing of Chem/Bio Dispersion Simulations," 2004, Proceedings of IASTED International Conference on Environmental Modeling and Simulation, St. Thomas, US Virgin Islands, 22-24 November 2004.

Bendat, J.S. and A.C. Piersol, 2000, *Random Data: Analysis and Measurement Procedures,* 3rd Edition, Wiley, New York.

Boris, J.P., 1971, "A Fluid Transport Algorithm That Works," International Center for Theoretical Physics, Trieste, Italy, 2-20 August 1970, published in *Computing as a Language of Physics*, International Atomic Energy Authority publication IAEA-SMR-9/18, 1971, pp. 171-189.

Boris, J.P. and D.L. Book, 1973, "Flux-Corrected Transport I. SHASTA, A Fluid Transport Algorithm That Works," *Journal of Computational Physics* **11(1)**, pp. 38-69, reprinted in the special anniversary edition of the *Journal of Computational Physics* **135(2)**, pp. 172-186, August 1997.

Boris, J.P. and D.L. Book, 1976, "Solution of Continuity Equations by the Method of Flux-Corrected Transport," Chapter 3 in *Methods in Computational Physics***, 16**, Academic Press, New York, pp. 85-129.

Boris, J.P., 1989, "On Large Eddy Simulation Using Subgrid Turbulence Models," *Whither Turbulence? Turbulence at the Crossroads: Lecture Notes in Physics* **357,** J.L. Lumley (ed), Springer–Verlag, New York, pp. 344–353.

Boris, J.P., F.F. Grinstein, E.S. Oran and R.L. Kolbe, 1992, "New Insights into Large Eddy Simulation," *Journal of Fluid Dynamics Research* **10 (4-6)**, pp. 199-228, (also U.S. Naval Research Laboratory Memorandum report NRL/MR/4400-92-6979).

Boris, J.P., A.M. Landsberg, E.S. Oran and J.H. Gardner, 1993, "LCPFCT - A Flux-Corrected Transport Algorithm for Solving Generalized Continuity Equations," U.S. Naval Research Laboratory Memorandum Report NRL/MR/6410-93-7192.

Boris, J.P., B.Z. Cybyk, and T.R. Young, 1999, "Computing Transport and Dispersion of Atmospheric Contaminants," Proceedings: 7th Annual Conference of the CFD Society of Canada, 1-2 June 1999, Halifax Nova Scotia, Canada.

Boris, J.P., 2001, "Reinventing Defense: Using High Performance Computing," Invited Lecture: DoD HPC User's Group Conference, Biloxi MS, 19-20 June 2001.

Boris, J.P., B.Z. Cybyk, T.R. Young, Jr., M.H. Emery and S.A. Cheatham, 2001a, "Simulation of Fluid Dynamics Around Complex Urban Geometries," *AIAA Paper 01-0803*, 39th Aerospace Sciences Meeting, Reno NV, 8 - 11 January 2001.

Boris, J.P., S.A. Cheatham, and T.R. Young, Jr., 2001b, "Simplifications Arising from Complexity in Urban Contaminant Transport," Annual meeting, APS Division of Computational Physics, Cambridge MA, 25-28 June 2001, *Bulletin of the American Physical Society* **46**(F), paper Q2.009, 25-28 June 2001.

Boris, J.P., 2002, "The Threat of Chemical and Biological Terrorism: Preparing a Response," *Computing in Science and Engineering* **4**(2), pp. 22–32.

Boris, J.P., K. Obenschain, G. Patnaik, and T. Young, Jr., 2002, "CT-Analyst: Fast and Accurate CBR Emergency Assessment," Proceedings: 1st Joint Conference on Battle Management for Nuclear, Chemical, Biological and Radiological Defense, Williamsburg VA, 4 – 8 November 2002.

Boris, J., S. Cheatham, G. Patnaik, M. McGinnis, K. Obernschain, T. Young, Jr. and T. Rappolt, 2004, "Variability, Validation, and CFD: There is No Truth!" Presentation and Proceedings, GMU Atmospheric Modeling Workshop, Fairfax VA, 14 July 2004.

Boris, J.P., C. Lind, K. Obenschain, G. Patnaik, and T. Young, Jr., 2005, "Faster and More Accurate CBR Emergency Assessment For Airborne Contaminants In Urban Environments," in *Urban Dispersion Modeling*, NATO Publication RTO-MP-AVT-120, presented at the NATO-AVT Workshop on Dispersion of Air Borne Matter in the Urban Environment, Brussels Belgium, 1-2 April 2004.

Britter, R.E. and Hanna, S.R., 2003, "Flow and Dispersion in Urban Areas," *Ann. Rev. Fluid Mech.* **35**, 2003, pp. 469-496.

Brown, M.J., J. Reisner, S. Smith and D. Langley, 2001, "High Fidelity Urban Scale Modeling, " *LA-UR-01-1422*, Los Alamos National Laboratory.

Calhoun, R., F. Gouveia, J. Shinn, S. Chan, D. Stevens, R. Lee, and J. Leone, 2004: Flow Around a Complex Building: Experimental and Large Eddy Simulation Comparisons, *Journal of Applied Meteorology* **(43)**, pp. 696-710.

Camelli, F., S. Hanna, and R. Lohner, 2004, "Simulation of the MUST Field Experiment Using the FEFLO-URBAN Model," 5th Urban Environment Conference, Vancouver.

Camelli F., R. Lohner and S. Hanna, 2005, "VLES Study of MUST Experiment," AIAA-2005-1279, 43rd AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada, 10-13 January 2005.

Castro, I. P. and A.G. Robins, 1977, "The flow around a surface-mounted cube in uniform and turbulent streams," *Journal of Fluid Mechanics* **79** part 2, pp. 307-335.

Chan, S, D. Stevens, and R. Lee, 2000, "A Model for Flow and Dispersion around Buildings and its Validation Using Laboratory Measurements," Third Symposium on the Urban Environment, Davis, CA, American Meteorological Society, pp. 56-57.

Chang J.C. and S.R. Hanna, 2004, "Air Quality Model Performance Evaluation," *Meteorology and Atmospheric Physics* (87), pp. 167-196.

Cybyk, B.Z., T.R. Young Jr., J.P. Boris, 1998, "Chapter 5 - Modeling and Simulation," *911-Bio Consequence Management ACTD Final Report,* U.S. Department of Defense Threat Reduction Agency, Nov 1998.

Fureby, C. and F.F. Grinstein, 2002, "Large Eddy Simulation of High Reynolds-Number Free and Wall Bounded Flows," *Journal of Computational Physics* **181**, pp. 68-86.

Gaydos, T.M. and I.J. Nyden, 2004, "Survey of Operationally Based Transport and Dispersion Software," Presentation and Proceedings, *Science and Technology for Chem-Bio Information Systems*, Williamsburg VA, 18-21October 2004.

Grinstein, F.F. and C. Fureby, 2004, "From Canonical to Complex Flows: Recent Progress on Monotonically Integrated LES," *Computational Science and Engineering* **6**, pp. 37-49.

Grinstein, F., L. Margolin and W. Rider (eds), 2006, *Implicit Large Eddy Simulation: Computing Turbulent Flow Dynamics*, Cambridge University Press, New York.

Hanna S.R., O.R. Hansen and S. Dharmavaram, 2004, "FLACS air quality CFD model performance evaluation with Kit Fox, MUST, Prairie Grass, and EMU observations. *Journal of the Atmospheric Environment* **(38)**, pp. 4675-4687.

Hendricks, E., Burrows, D.A., Diehl, S., and Keith, R., 2004, "Dispersion in the Downtown Oklahoma City Domain: Comparisons Between the Joint Urban 2003 Data and the RUSTIC/MESO Models," in *5th AMS Symposium on the Urban Environment*, Vancouver.

Huber, A., M. Freeman and K. Kuehlert, 2003, "Environmental CFD simulation and visualization: Examples in Support of the Reconstruction of the Smoke/Dust Plume from the World Trade Center Site Following the Events of September 11, 2001," Presented at: Science Forum 2003, Washington, DC, 5-7 May 2003.

Hunt, J.C.R., C.J. Abell, J.A. Peterka and H. Woo, 1978, "Kinematical studies of the flows around free or surface-mounted obstacles; applying topology to flow visualization," *Journal of Fluid Mechanics*, **86** part 1, pp. 179-200.

Ikeda, T. and T. Nakagawa, 1979, "On the SHASTA FCT Algorithm for the Equation $d\rho/dt + d(\rho V)/dx = 0$," *Mathematics of Computation* **33** (148), pp. 157-1169.

Iselin, J.P., G. Patnaik, C. Cox, J.P. Boris, 2002, "Effect of Inflow Boundary Parameters on Boundary Layer Transition," *Proceedings: Fourth George Mason University Transport and Dispersion Modeling Workshop*, George Mason University, Fairfax VA. July 11-12, 2002,

Koomullil, R.P., and B.K. Soni, 2001, "Wind Field Simulations in Urban Area," AIAA Paper AIAA-2001-2621. *15th AIAA Computational Fluid Dynamics Conference*, Anaheim CA, 11-14 June 2001.

Kuzmin, D., R. Lohner and E. Turek (eds), 2005, *Flux-Corrected Transport*, Spring-Verlag, Berlin Heidelberg.

Lakehal, D. and W. Rodi, 1997, "Calculation of the flow past a surface-mounted cube with two-layer turbulence models," *Journal of Wind Eng. and Ind. Aerodynamics*, **67&68**, pp. 65-78.

Landsberg, A.M., J.P. Boris, W.C. Sandberg, and T.R. Young Jr**.**, 1995, "Analysis of the Nonlinear Coupling Effects of a Helicopter Downwash with an Unsteady Ship Airwake," AIAA Paper 95-0047, 33rd Aerospace Sciences Meeting, Reno NV, 9-12 January 1995, AIAA, Reston VA.

Landsberg, A.M. and W.C. Sandberg, 2000, "DDG-51 FLT-IIA Airwake Study Part 3: Temperature Field Analysis for Baseline and Upgrade Configurations," U.S. Naval Research Memorandum Report, NRL/MR/6401-00-8432.

Lien, F. S. and E. Yee, 2004, "Numerical Modeling of the Turbulent Flow Developing Within and Over a 3-D Building Array, Part 1: A High-Resolution Reynolds-Averaged Navier-Stokes Approach," *Boundary-Layer Meteorology* **112,** pp. 427–466.

Lind, C.A., J.P. Boris, E.S. Oran, W.J. Mitchell, and J.L. Wilcox, 1997, "The Response of an Open Air Detonation Facility to Blast Loading," *Structures Under Extreme Loading Conditions: ASME PVP* - Vol. 351, pp. 109-126.

Martinuzzi, R. and Tropea, C., 1993, "Flow Around Surface-Mounted, Prismatic Obstacles Placed in a Fully Developed Channel Flow," *Journal of Fluids Engineering*, **115**, pp. 85-92.

Obenschain, K., G. Patnaik, and J.P. Boris, 2004, "Using CT-ANALYST to Optimize Sensor Placement," *Proceedings: Chemical and Biological Sensing* **V (5416),** *SPIE Defense and Simulation Symposium*, P. Gardner (ed), SPIE Paper 5416-02, Orlando FL, 12-13 April 2004, pp. 14-20.

Oberkampf, W. and J. Helton, 2002, "Investigation of Evidence Theory for Engineering Applications," AIAA Paper AIAA-2002-1569, 43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, Denver, Colorado, Apr. 22-25, 2002

Oberkampf, W. and Barone, 2004, "Measures of Agreement Between Computation and Experiment: Validation Metrics," AIAA Paper, AIAA-2004-2626, 34th AIAA Fluid Dynamics Conference and Exhibit, Portland, Oregon, June 28-1, 2004.

Oran, E.S. and J.P. Boris, 1993, "Computing Turbulent Shear Flows - A Convenient Conspiracy", *Computers in Physics* **7(5),** pp. 523-533.

Oran, E.S. and J.P. Boris, 2001, *Numerical Simulation of Reactive Flow* (2nd edition), Cambridge University Press, New York.

Patnaik, G., J.P. Boris, F.F. Grinstein, and J.P. Iselin, 2003, "Large Scale Urban Simulations with the MILES Approach," *AIAA Paper 2003-4104,* AIAA CFD Conference, Orlando FL, 25 June 2003.

Patnaik, G., J. Boris, F. Grinstein, and J**.** Iselin, 2005, "Large Scale Urban Simulations with FCT," in *Flux-Corrected Transport: Principles, Applications and Algorithms*, D. Kuzmin, R Lohner, and S. Turek (eds), 30th Anniversary FCT Conference, Dortmund Germany, 28 Sept-1 Oct 2003, (Springer-Verlag, Berlin-Heidelberg, 2005), pp. 105-130.

Patnaik, G., F. Grinstein, J. Boris, T. Young, Jr., and O. Parmhed, 2006, "Large Scale Urban Simulations," in *Implicit Large Eddy Simulation: Computing Turbulent Flow Dynamics*, F. Grinstein, L. Margolin, and W. Rider (eds), Cambridge University Press, New York, (in press).

Pullen, J., J. Boris, T. Young, G. Patnaik and J. Iselin, 2005, "A Comparison of Contaminant Plume Statistics from a Gaussian Puff and an Urban CFD Model for Two Large Cities," *Journal of the Atmospheric Environment* **(39),** pp. 1049-1068.

Rappolt, T.J., 2002, "Measurements of Atmospheric Dispersion in the Los Angeles Urban Environment: Summer 2001," *Tracer ES&T, Inc. Project Report 1322*, 28 February 2002.

Sagaut, P., 2004, *Large Eddy Simulation for Incompressible Flows*, 2nd Edition, Springer-Verlag.

Smith, W.S., M. Brown and D. DeCroix, 2002, "Evaluation of CFD Simulations Using Laboratory Data and Urban Field Experiments, " in *4th AMS Symposium on the Urban Environment,* Norfolk VA.

Snyder, W.H., 1994, "Some observations of the influence of stratification in building wakes," in *Stably Stratified Flows: Flow and Dispersion over Topography*, I.P. Castro and N.J. Rockliff, eds. Clarendon Press, pp. 301-324.

Stefaniw, E., T.R. Young,Jr. and J.P. Boris, 1998, "FAST3D: A High Fidelity Transport and Diffusion Model," *Chem-Bio/Smoke Modeling and Simulation Newsletter* **(4)** 1, Chemical and Biological Defense Information Analysis Center. Aberdeen Proving Ground, MD. January 1998.

Warner, S., N. Platt, J. Heagy, S. Bradley, G. Bieberbach, G. Sugiyama, J. Nasstrom, K. Foster and D. Larson, 2001, User-Oriented Measures of Effectiveness for the Evaluation of Transport and Dispersion Models," IDA Paper P-3554.

Wirsching, P., T. Paez, and K. Ortiz, 1995, *Random Vibrations Theory and Practice*, Wiley, New York.

Yellen, L. and W. Jacobson, 2003, "Fox News Exclusive," Fox Evening News WFLD-TV/WPWR-TV, first aired 31 January 2003.

Young, T.R. Jr., A.M. Landsberg and J.P. Boris, 1993, "Implementation of the Full 3D FAST3D (FCT) Code Including Complex Geometry on the Intel iPSC/860 Parallel Computer," *High Performance Computing 1993 -- Grand Challenges in Computer Simulation*, Adrian Tentner (ed), 1993 SCS Simulation Multiconference, Arlington VA, 29 March - 1 April 1993, SCS, San Diego.

Young, T., J. Boris, S. Hooper, C. Lind, K. Obenschain, and G. Patnaik, 2004, "Emergency Assessment With Sensors And Buildings: Advances In CT-Analyst Technology," Presentation and Proceedings, *Science and Technology for Chem-Bio Information Systems*, Williamsburg VA, 18-21October 2004.

Figures

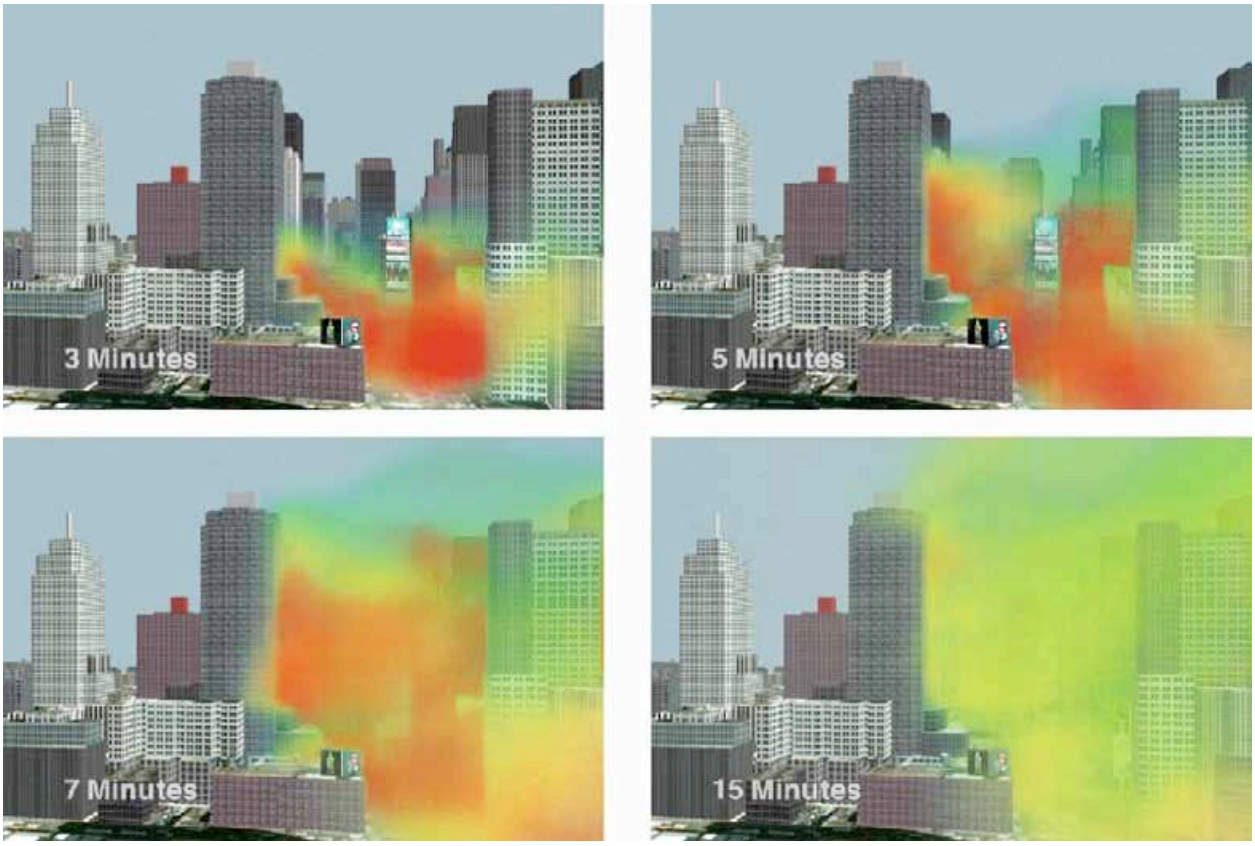

Figure 1. Three-dimensional visualizations of a FAST3D-CT CFD simulation. The source was an instantaneous release of a neutrally buoyant tracer gas at ground level in Times Square, New York City. The frames show relative tracer concentrations at 3, 5, 7, and 15 minutes after release. (Graphics: Robert Doyle)

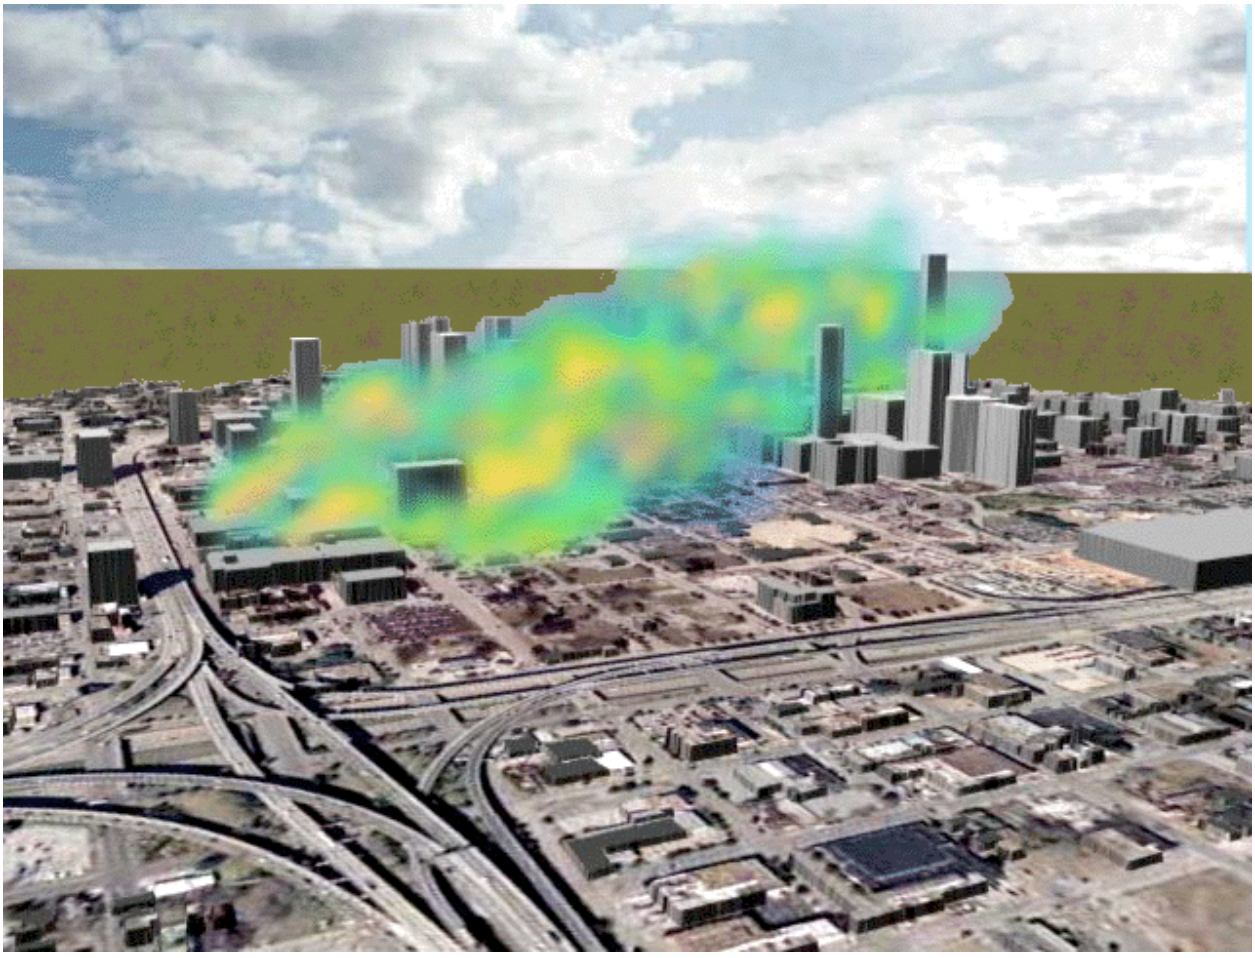

Figure 2. Contaminant dispersion from an instantaneous release in Houston Texas. The cloud is visualized 30 minutes after release. Orange and yellow regions are the highest concentrations. Note the rapid vertical mixing controlled by the buildings. This is one of the validation scenarios shown in Figure 19. (Graphics: Robert Doyle)

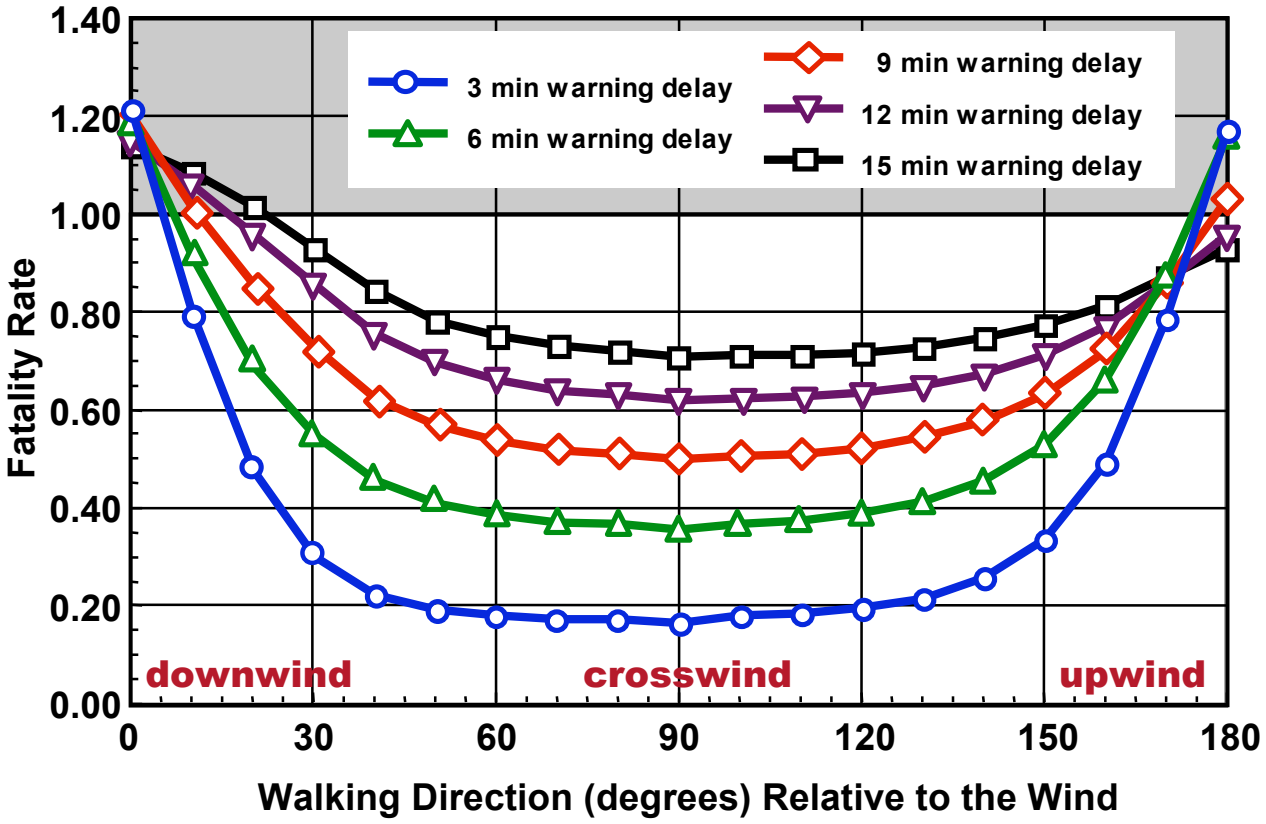

Figure 3. The dilemma for urban civil defense from airborne contaminants. The cost in lives of delay can be very high. Walking perpendicular to the wind allows most of the people to escape lethal exposures if the warning is issued quickly enough and accompanied by the correct direction to walk.

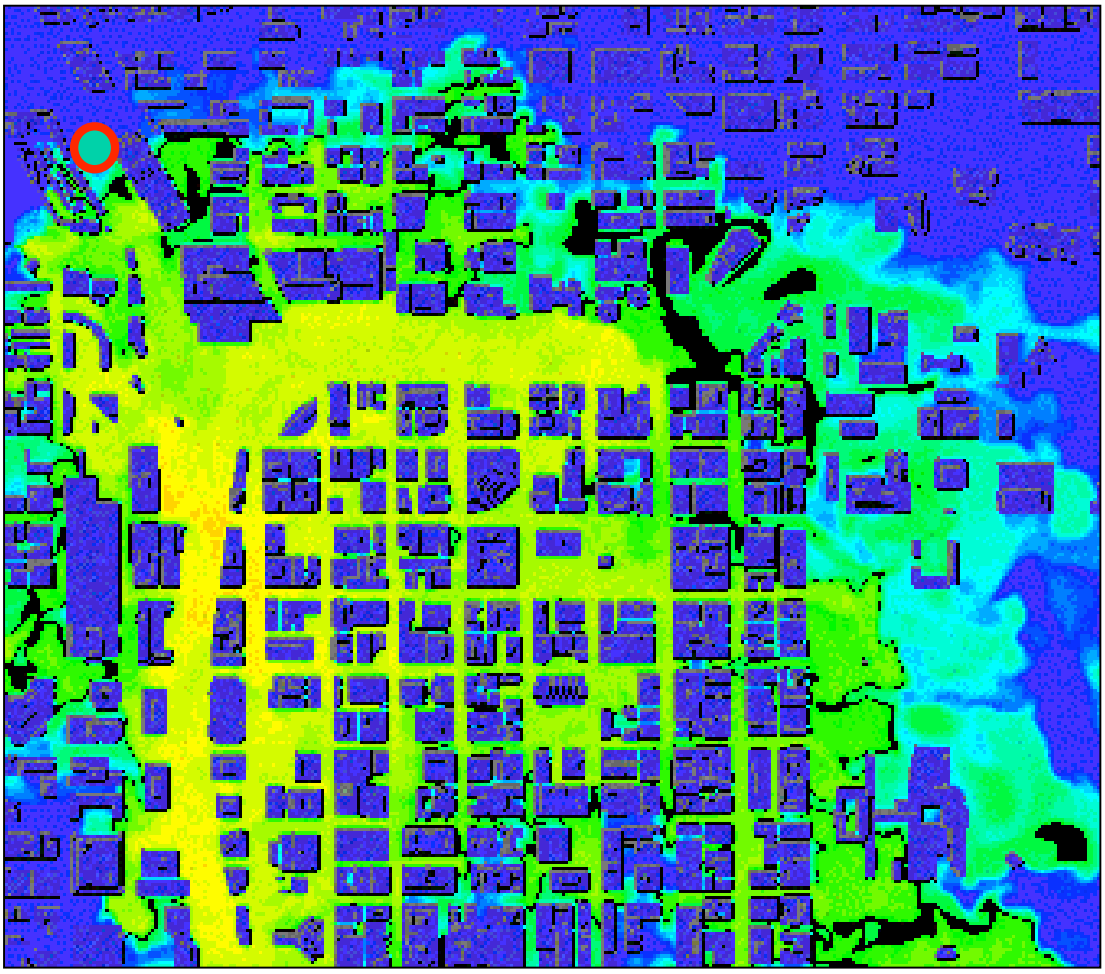

Figure 4. Cross-section of a contaminant cloud at ground level computed by FAST3D-CT for downtown Chicago. The prevailing wind is from 320 degrees at 3 m/s and the visualization is 18 minutes after release of the contaminant. The source location and wind conditions correspond exactly to the CT-Analyst scenario shown below in Figure 13. The source was an instantaneous (explosive) release in the red circle.

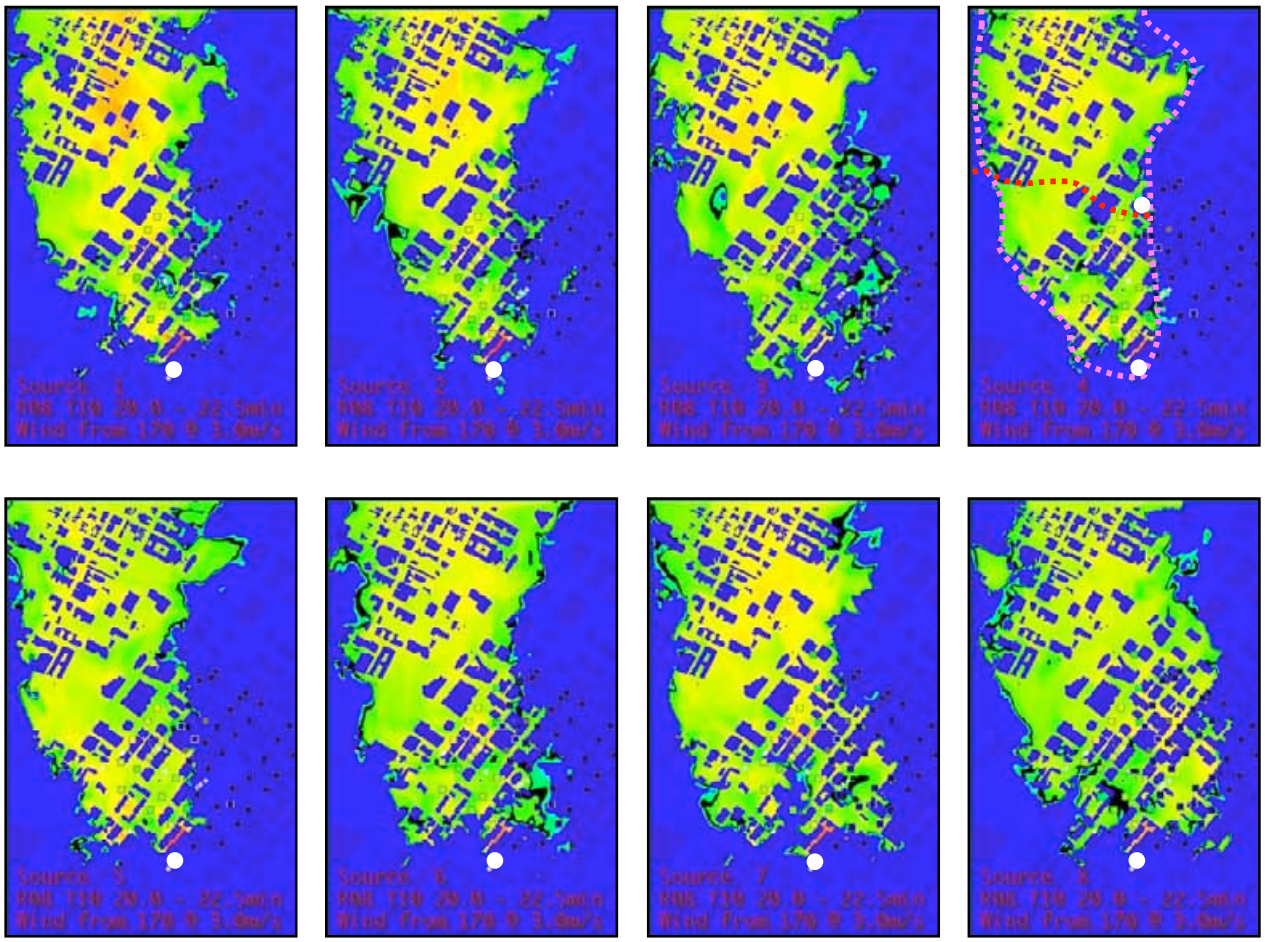

Figure 5. Eight independent realizations of a single source are compared with Tracer ES&T sampler data taken in downtown Los Angeles (colored squares). The wind is from 170^0 at 3 m/s with moderate wind fluctuations. Release times separated by five minutes cause the appreciable differences shown as the wind fluctuations interact with the complex building vortex shedding. The source is the white circle in the lower center of each panel.

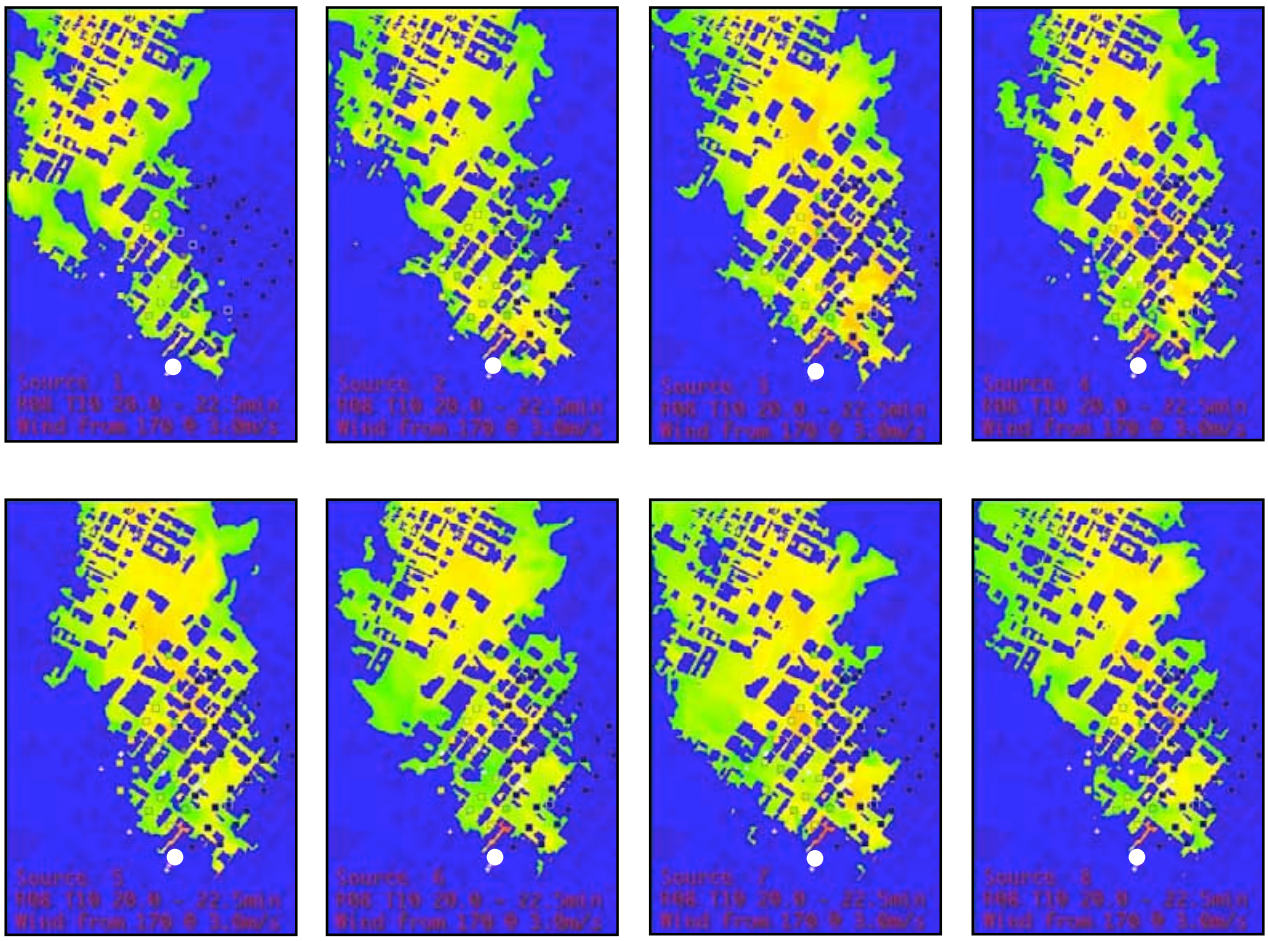

Figure 6. Eight independent realizations of a single source computed by FAST3D-CT as in Figure 5. The impressed urban boundary layer profile in this case has no wind fluctuations. Building vortex shedding and atmospheric instability account for all of the variability. The wind was from 170^0 at 3 m/s. Note that these concentrations appear a little larger near the source (more orange) than in Figure 5. The source is the white circle in the lower center of each panel.

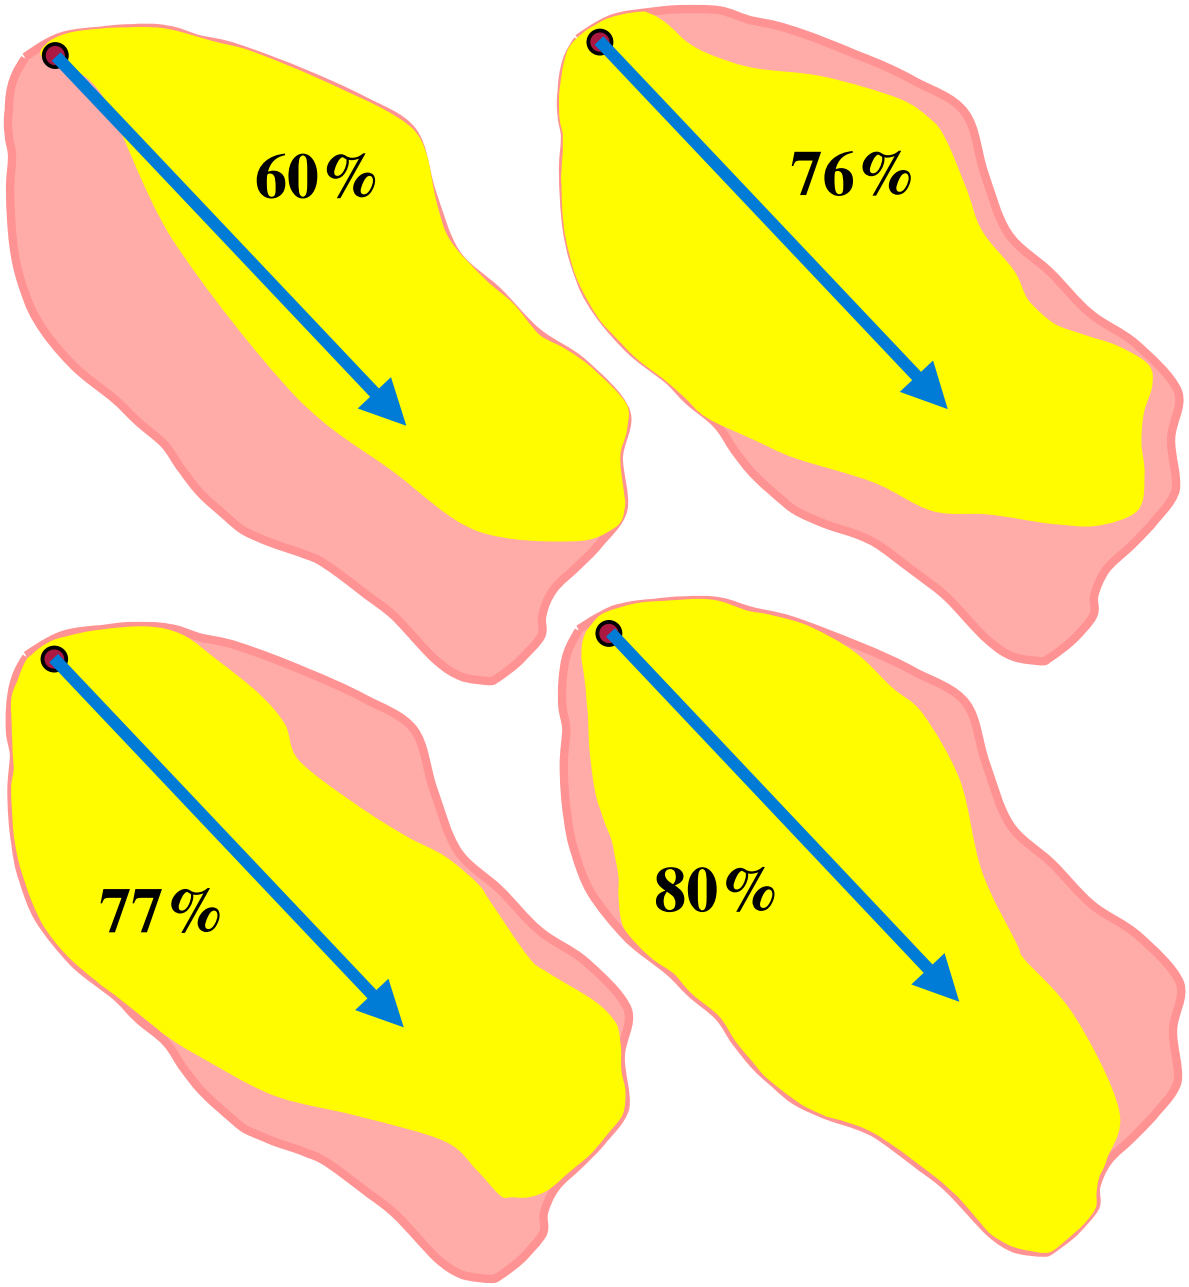

Figure 7. Diagram showing how a plume envelope (pink) is defined as the union of the areas with concentrations above a threshold value for a number of independent plume realizations (indicated by yellow above). Each of the four realizations shown fills a different fraction of the plume envelope.

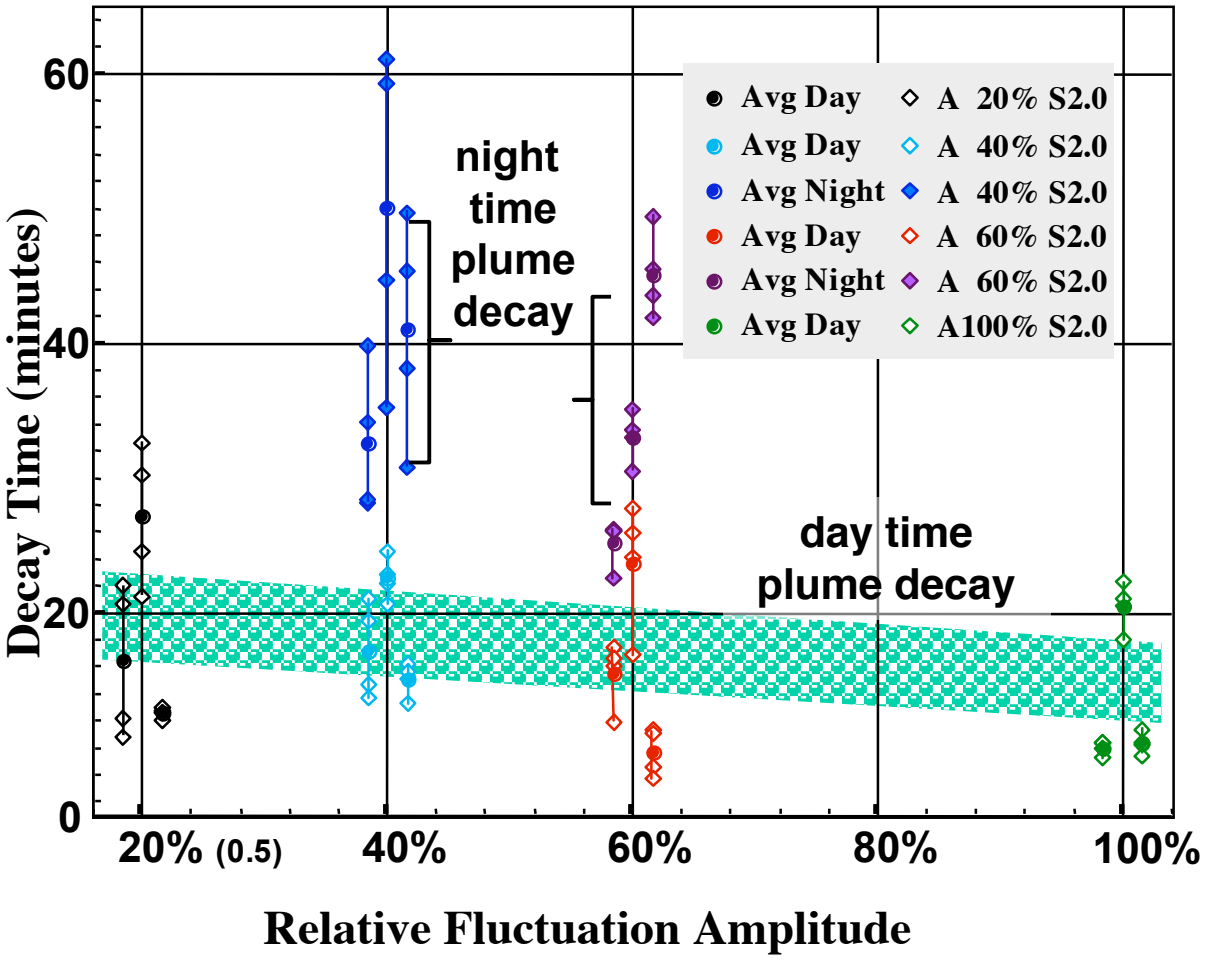

Figure 8. Contaminant decay times in an urban area depend on the strength of the wind fluctuations and on the stability of the atmosphere (time of day). The time of day has a bigger effect than the impressed fluctuation strength because the fluctuations caused by building vortex shedding are the dominant effect.

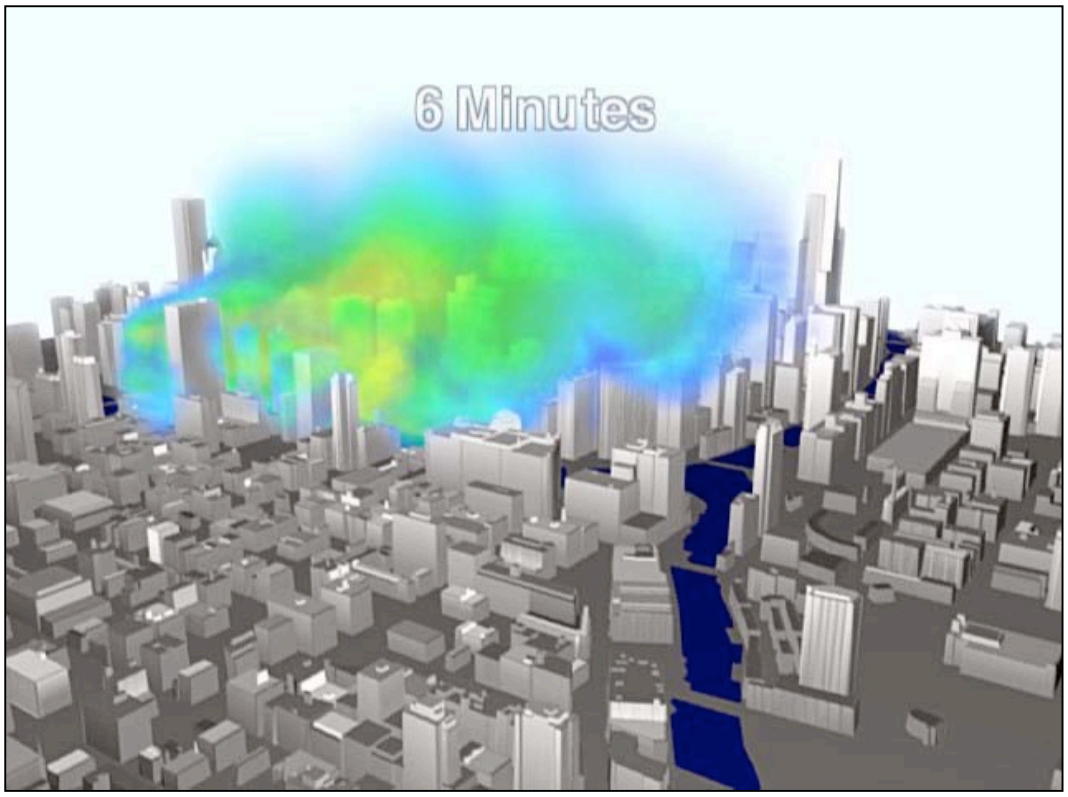

Figure 9. Visualization of a FAST3D-CT contaminant transport simulation over downtown Chicago with wind from the east at 3 m/s. The fountain effect is fully formed 6 minutes after the contaminant is released and is clearly visible behind the tall building in the upper left of the figure. (Graphics: Robert Doyle)

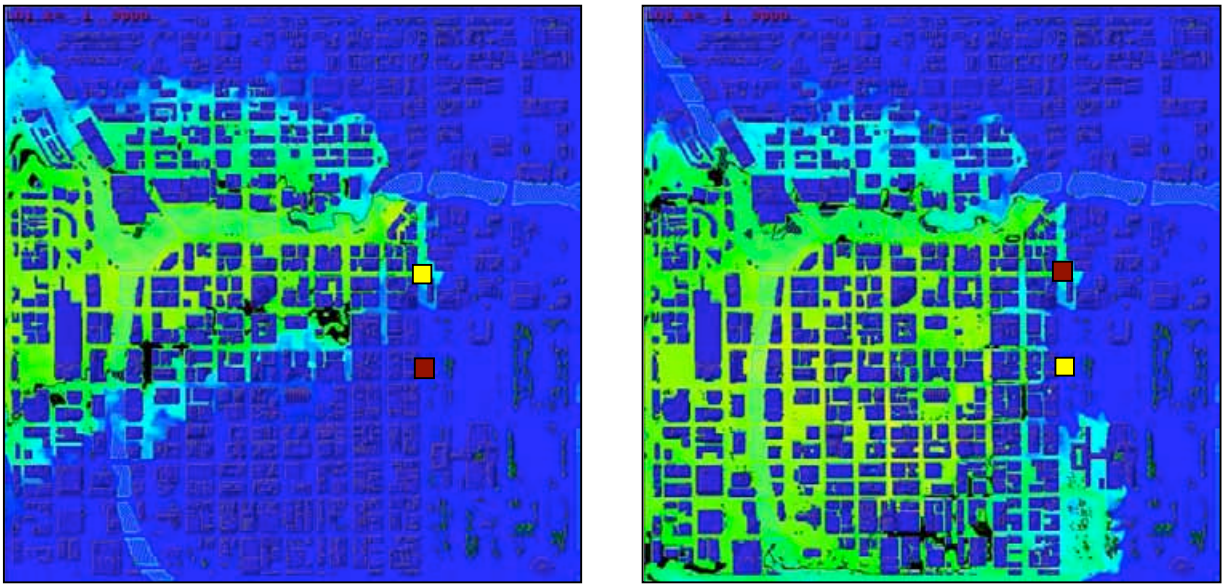

Figure 10. Two short-duration sources released 350 meters apart simultaneously in downtown Chicago. The urban geometry alone accounts for the differences.

Figure 11. Imagery (left) and a color-coded land use approximation (right) to the three-dimensional geometry of an urban region. Pink areas on the left indicate suburban developments, tan areas indicate low urban areas, and lavender and purple indicate moderate-building and high-building urban areas respectively.

Figure 12. Digital Elevation Map (DEM) color coded for land use as prepared automatically from lidar data. Right - enlargement of the indicated region on the left panel. Heights can be accurate to a fraction of a meter and horizontal resolution is one meter.

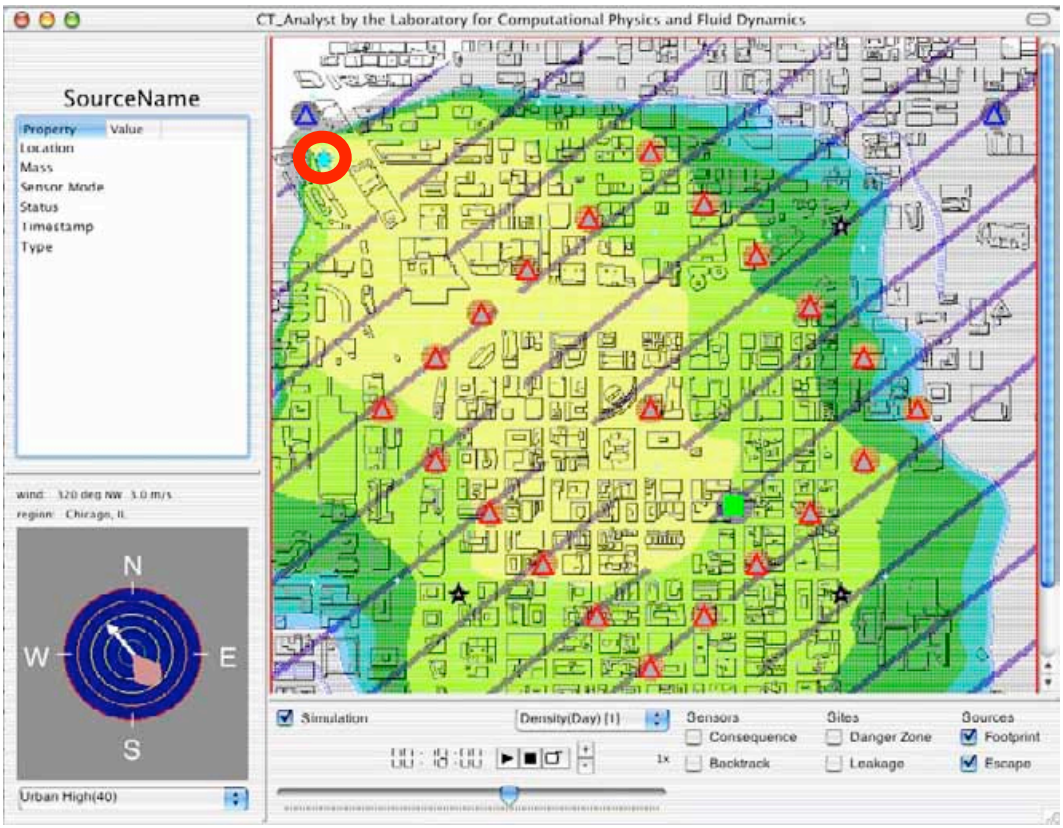

Figure 13. CT-Analyst full screen display showing contaminant concentration contours (yellow, green, and blue), contamination footprint (grey), and evacuation routes (magenta/purple) overlaid on a city map of downtown Chicago.

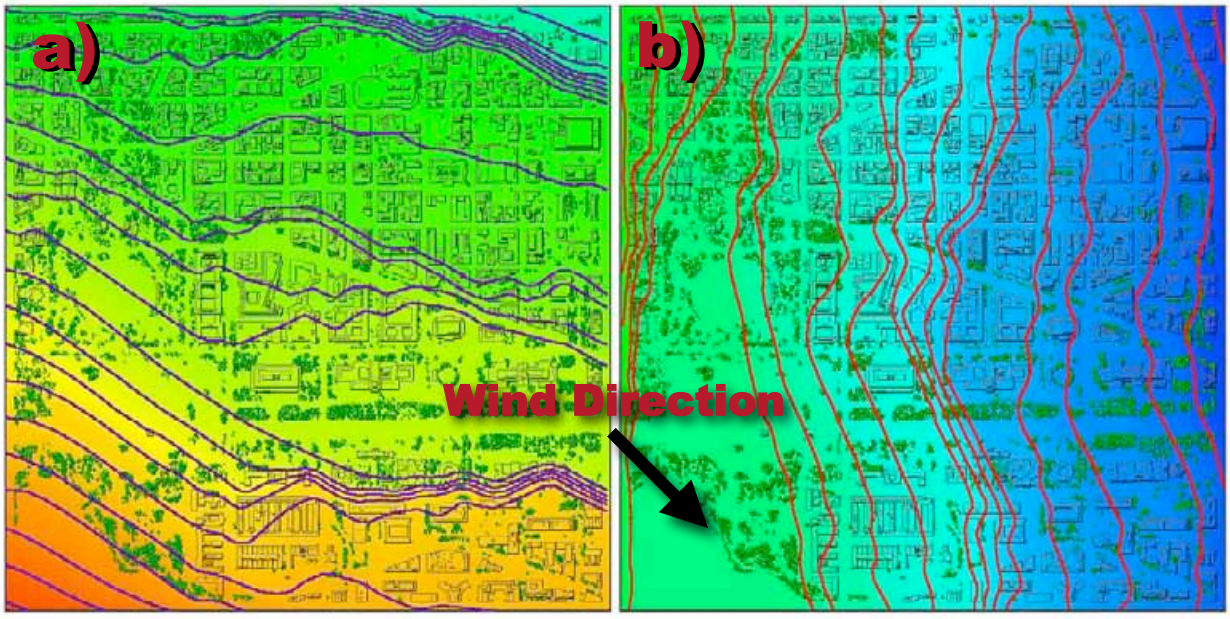

Figure 14. Cloud edge arrays used to construct a dispersion nomograf for a wind from the northwest. The left (a) and right (b) cloud edge arrays are real arrays that define nested, monotone sets of contours which are the limiting lines for the edges of an expanding contaminant cloud. The cloud edge arrays encode the building geometry effects.

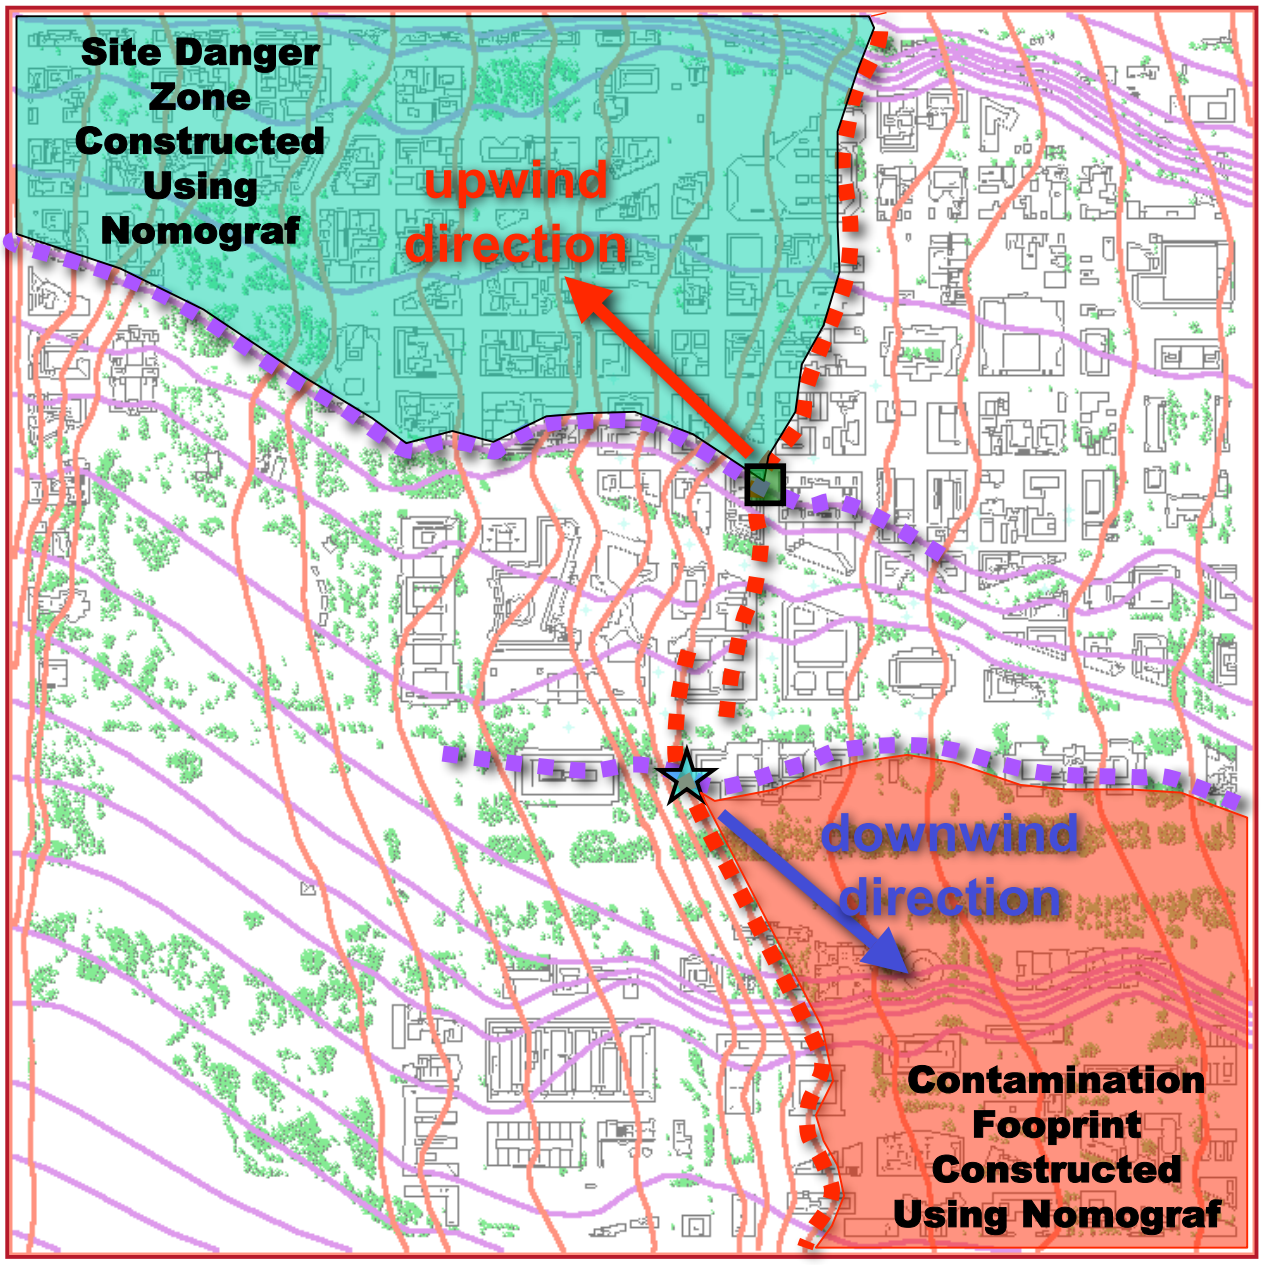

Figure 15. Overlapping the contours of left and right cloud edge arrays constitutes a dispersion nomograf™ for the computed wind direction. The cloud edge contours from Figure 14 are used to define upwind sectors (site danger zones and sensor backtrack regions) and downwind sectors (contaminant footprints and sensor consequence regions).

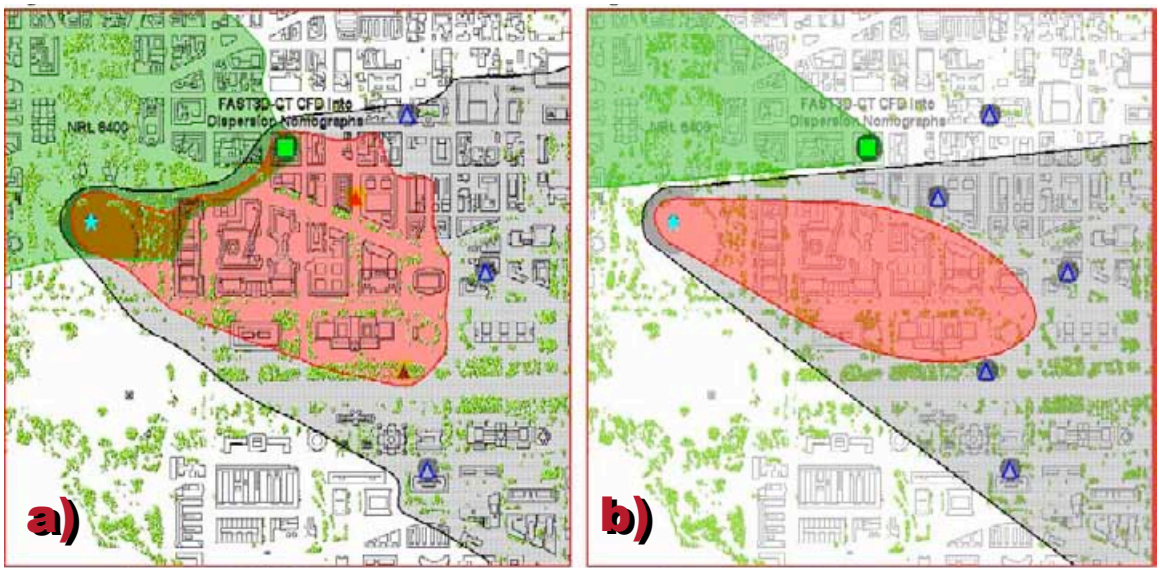

Figure 16. CT-Analyst nomografs capture building aerodynamics. a) Nomograf results with full building encoded. b) Nomograf computed excluding buildings, trees, and terrain. The figure shows the corresponding contamination footprints (gray), plume envelopes after 6 minutes (pink), and the upwind danger zone of a site (green).

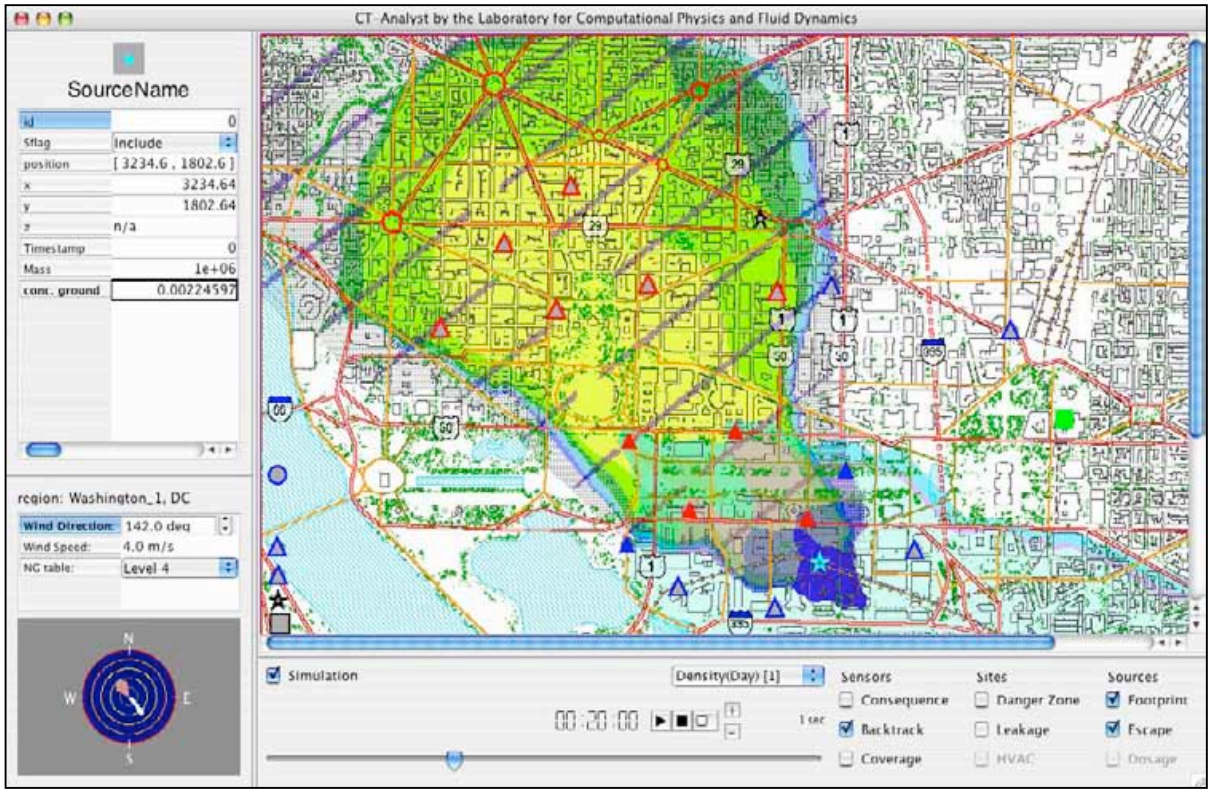

Figure 17. CT-Analyst full-screen display from a moderate area urban nomograf showing contaminant concentration contours (yellow, green, and blue), the contamination footprint (grey), and evacuation routes (purple lines) overlaid on a city map. The wind is from the south-east. Evacuation routes are optimized to minimize inhaled contaminant doses.

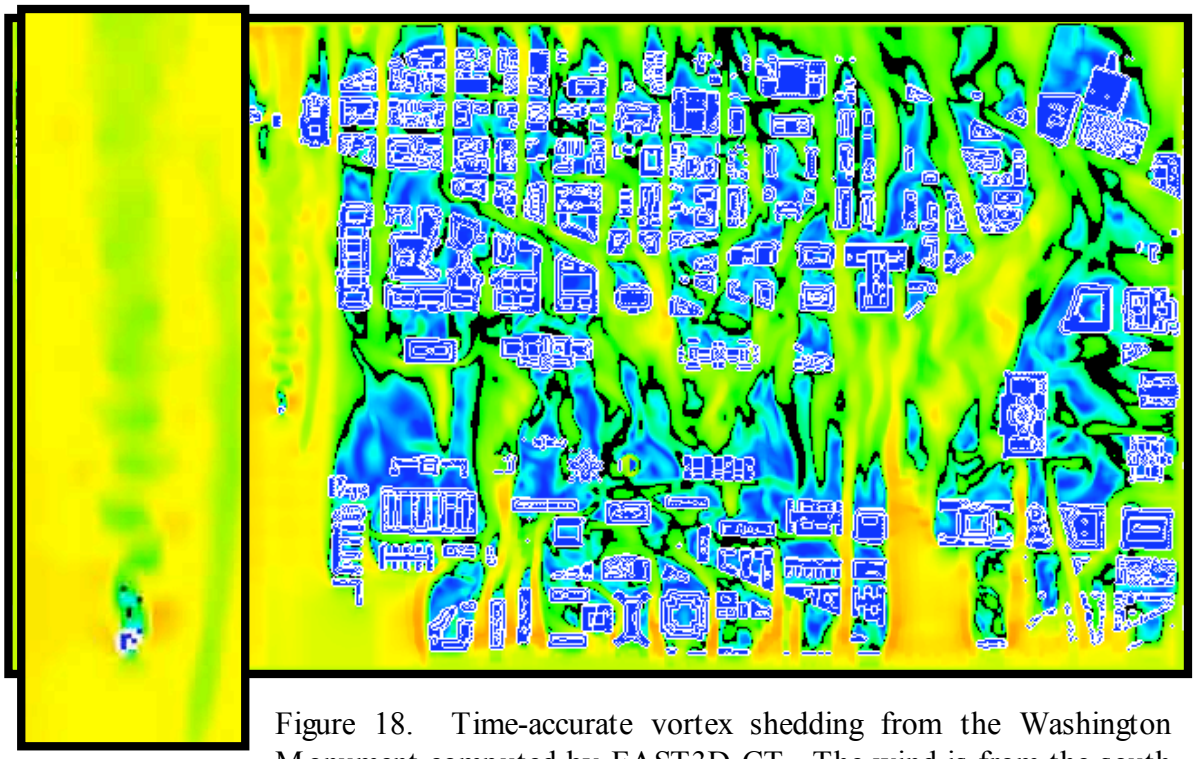

Figure 18. Time-accurate vortex shedding from the Washington Monument computed by FAST3D-CT. The wind is from the south (bottom) at 3 m/s. Vy is visualized 20m above the ground level.

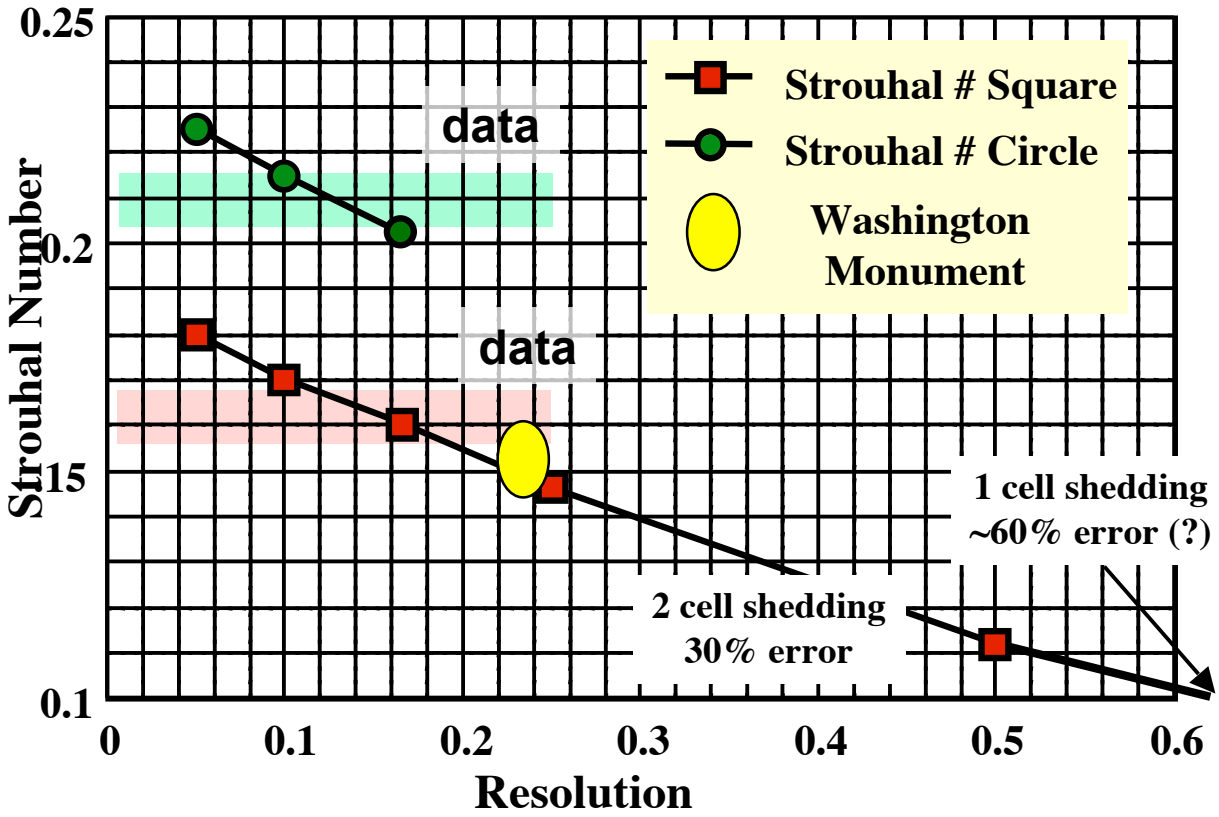

Figure 19. Strouhal number for vortex shedding computed by FAST3D-CT. The error is 3 to 5 % at the 5-meter resolution used in the detailed simulation of the Washington Monument.

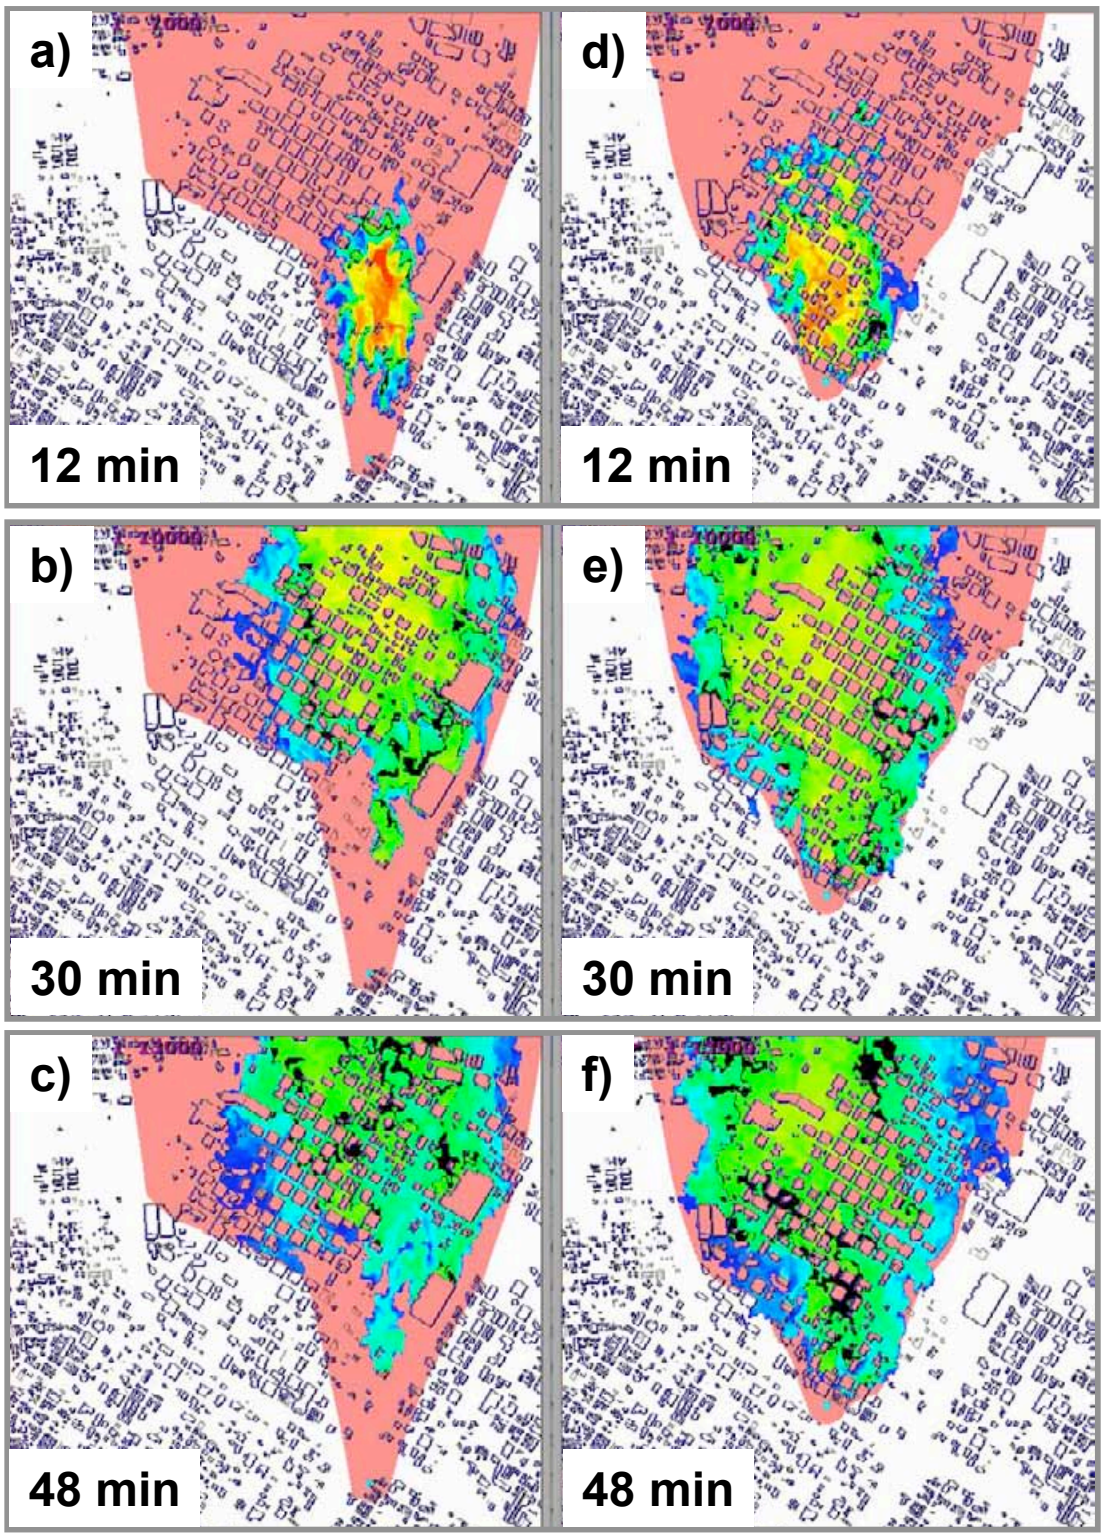

Figure 20. Evolution of FAST3D-CT concentration plots overlaid on the corresponding CT-Analyst contamination footprint (pink) area. CT-Analyst predictions are conservative by design but the influence of buildings and the close correspondence between the predictions is evident. Panels a, b, and c are for Source 1. d, e, and f are for Source 2.

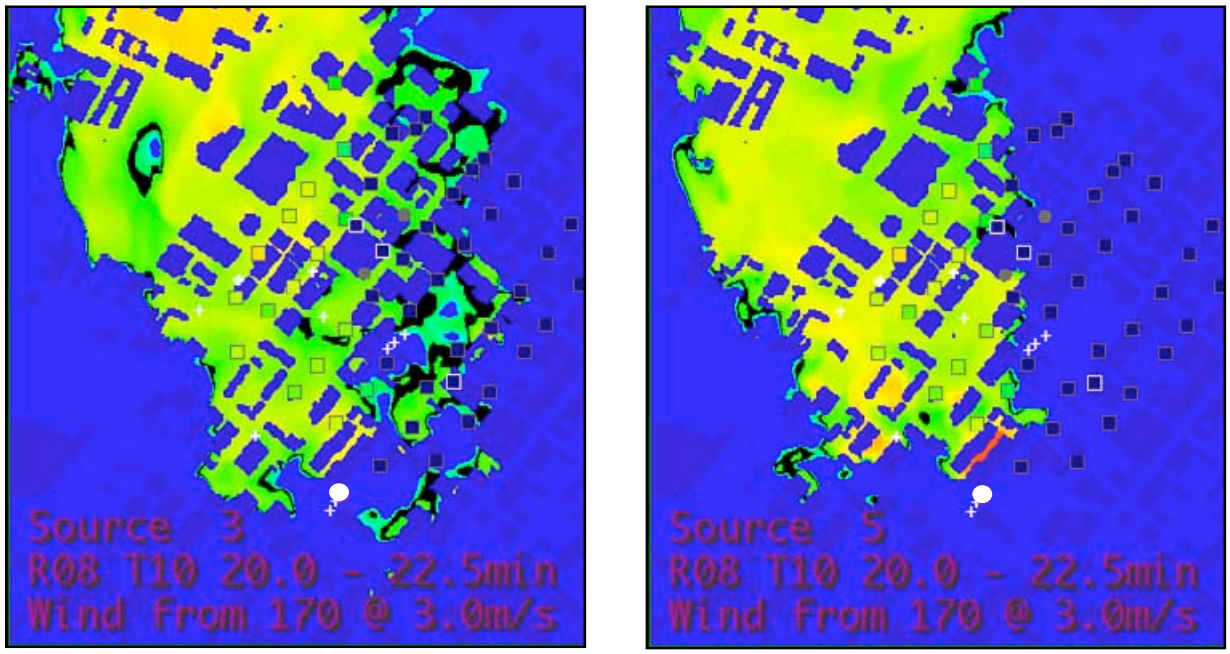

Figure 21. Enlargement of two of the eight FAST3D-CT realizations from Figure 5 above. Squares contain the Tracer ES&T, Inc. sampler data from Los Angeles release #8 plotted on the same color scale as the simulations. The source location is the white circle in each panel.

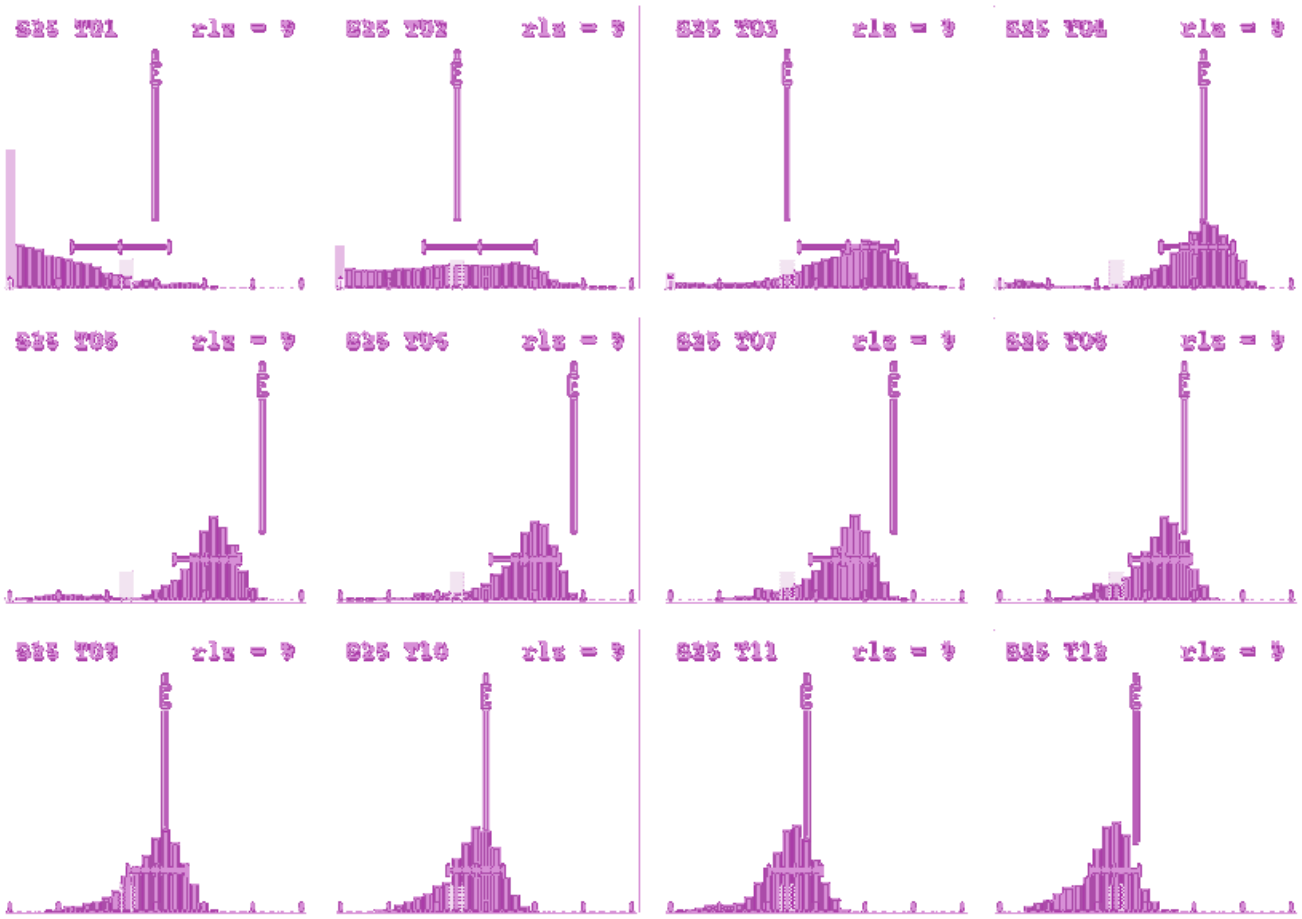

Figure 22. Composite concentration distribution functions from eight FAST3D-CT realizations at the location of Sampler 25 in the baseline FAST3D-CT simulation for Los Angeles Trial #8. The vertical green lines indicate the experimental value and the short blue cross-hatched bar indicates the assumed experimental threshold value of 20 parts per trillion used for this comparison. The locally prevailing wind was 3 m/s from 170^0 with modest wind fluctuations. The Chi-square value for this set of twelve measurements is ~13 and the corresponding probability of agreement is 45% (Abramowitz and Stegun, 1964).

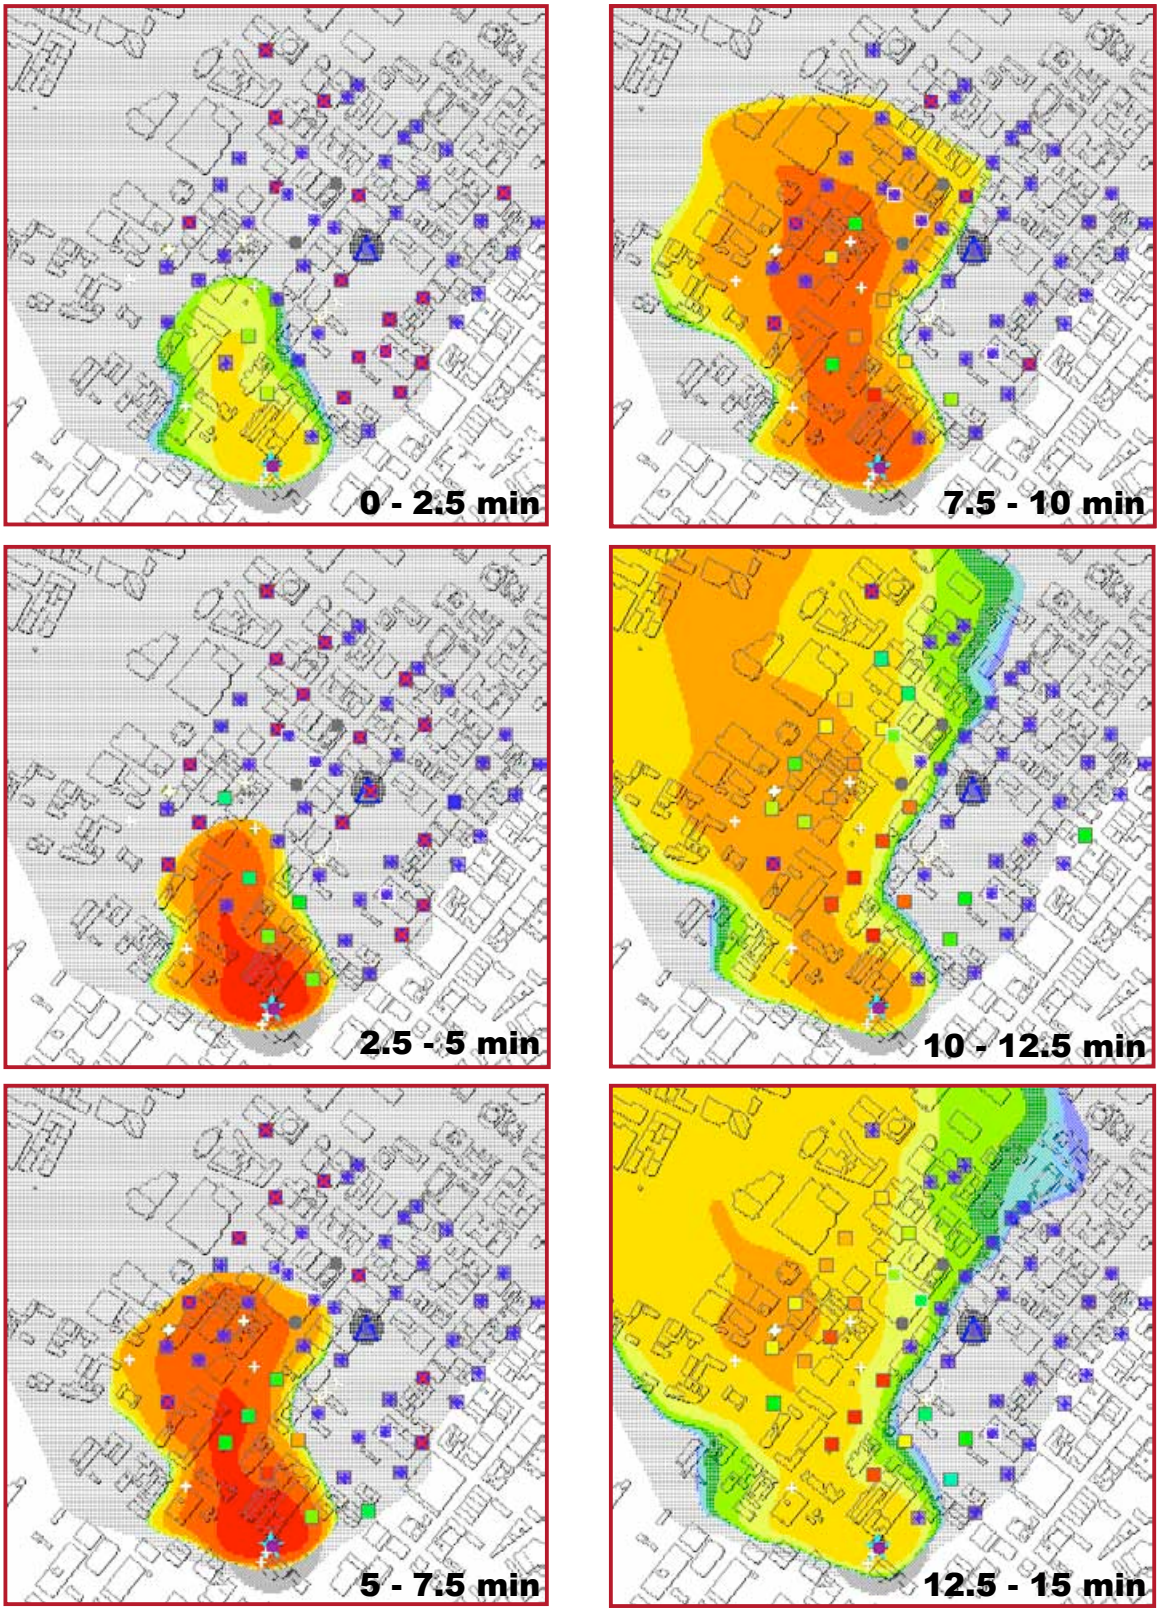

Figure 23a. Tracer ES&T sampler data (colored squares) on the corresponding CT-Analyst predictions for a simulated 5 minute SF6 release in Los Angeles. The first six sampling intervals are shown. The color tables are not identical.

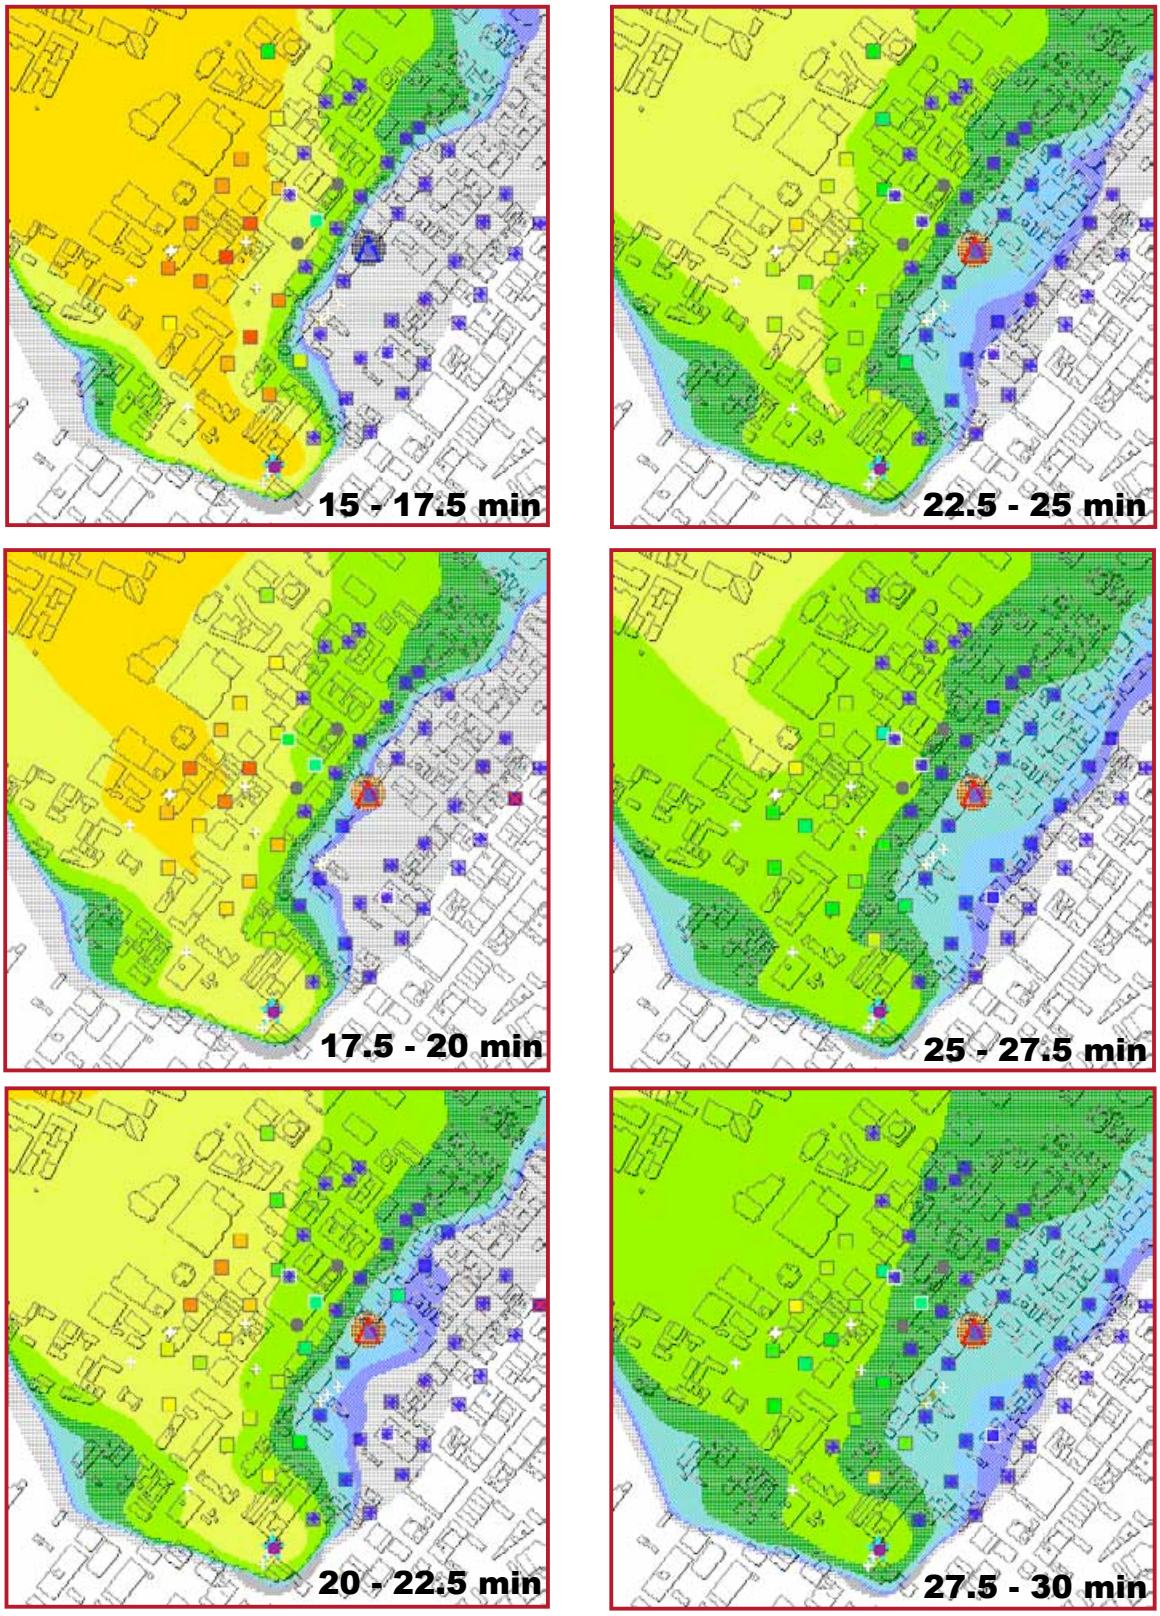

Figure 23b. Tracer ES&T sampler data (colored squares) on the corresponding CT-Analyst predictions for a simulated 5 minute SF6 release in Los Angeles. The last six sampling intervals shown. The gray area is the eventual contamination footprint.

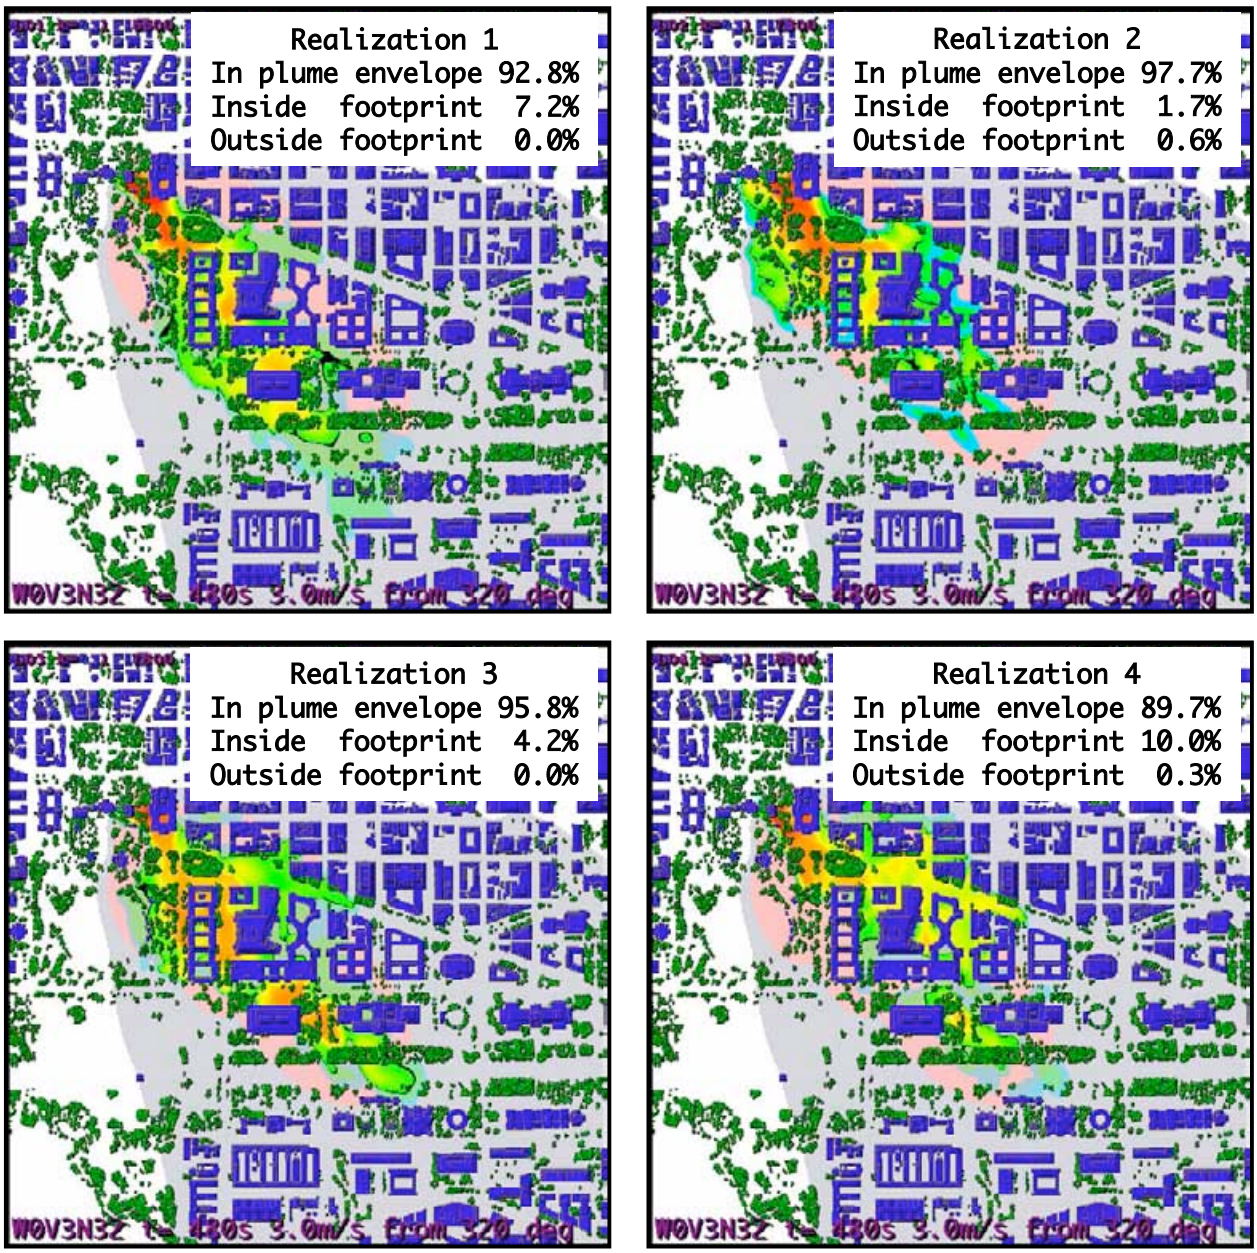

Figure 3.4. Four realizations of a scenario, computed using FAST3D-CT, were superimposed on the corresponding CT-Analyst plume envelope 8 minutes after the contaminant was released. The realizations differ in the phases of imposed wind fluctuations. The plume envelope is designed to be generally conservative. The Cumulative Figure of Merit (CFoM) is 0.853 at this time. The CF0M is shown as a function of time in Figure 3.5 for realizations 1 - 4 of Run N320.

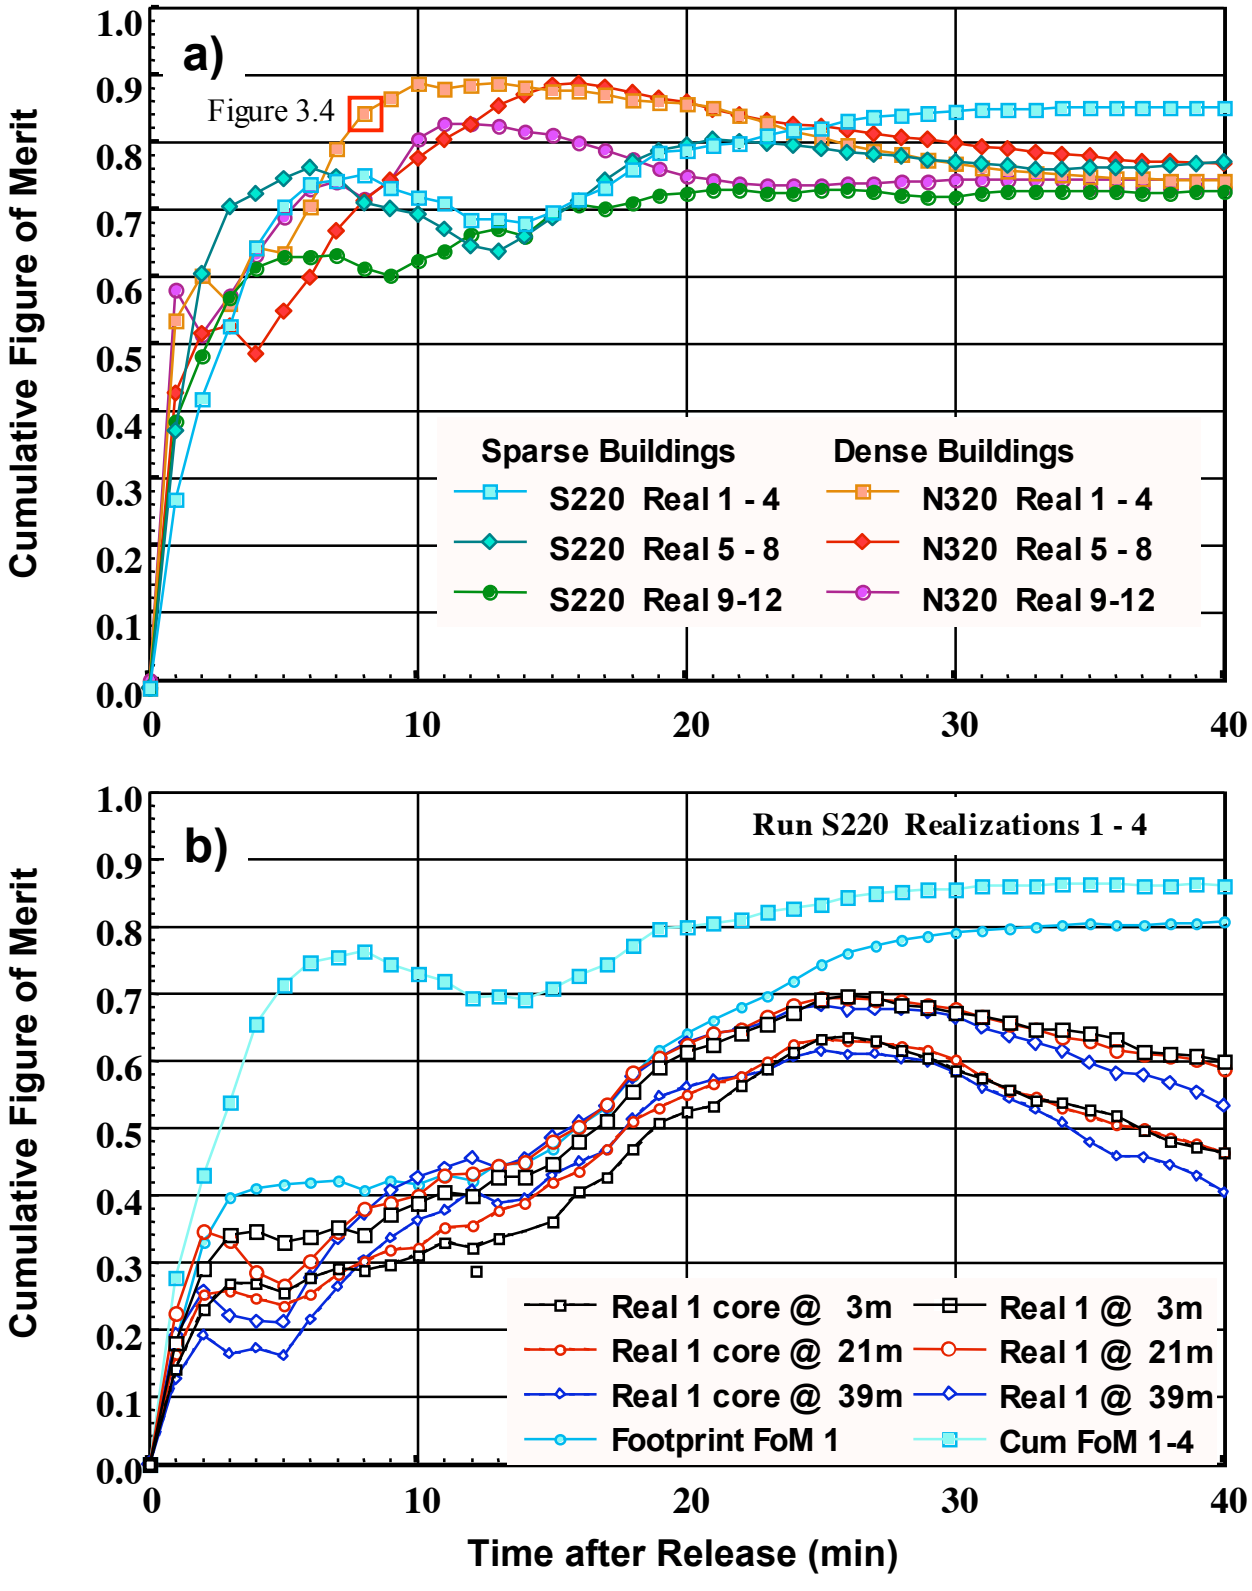

Figure 3.5 a) Nomograf Cumulative Figure of Merit for four realizations of sources at each of six locations. Run N320 corresponds to data from Figure 3.4. b) Turbulent transport from wind gusts and buildings limits vertical variations within the urban canopy. Sparse buildings and strong wind fluctuations characterize Run S220.

FAST3D-CT Simulations Converge with Increasing Grid Resolution

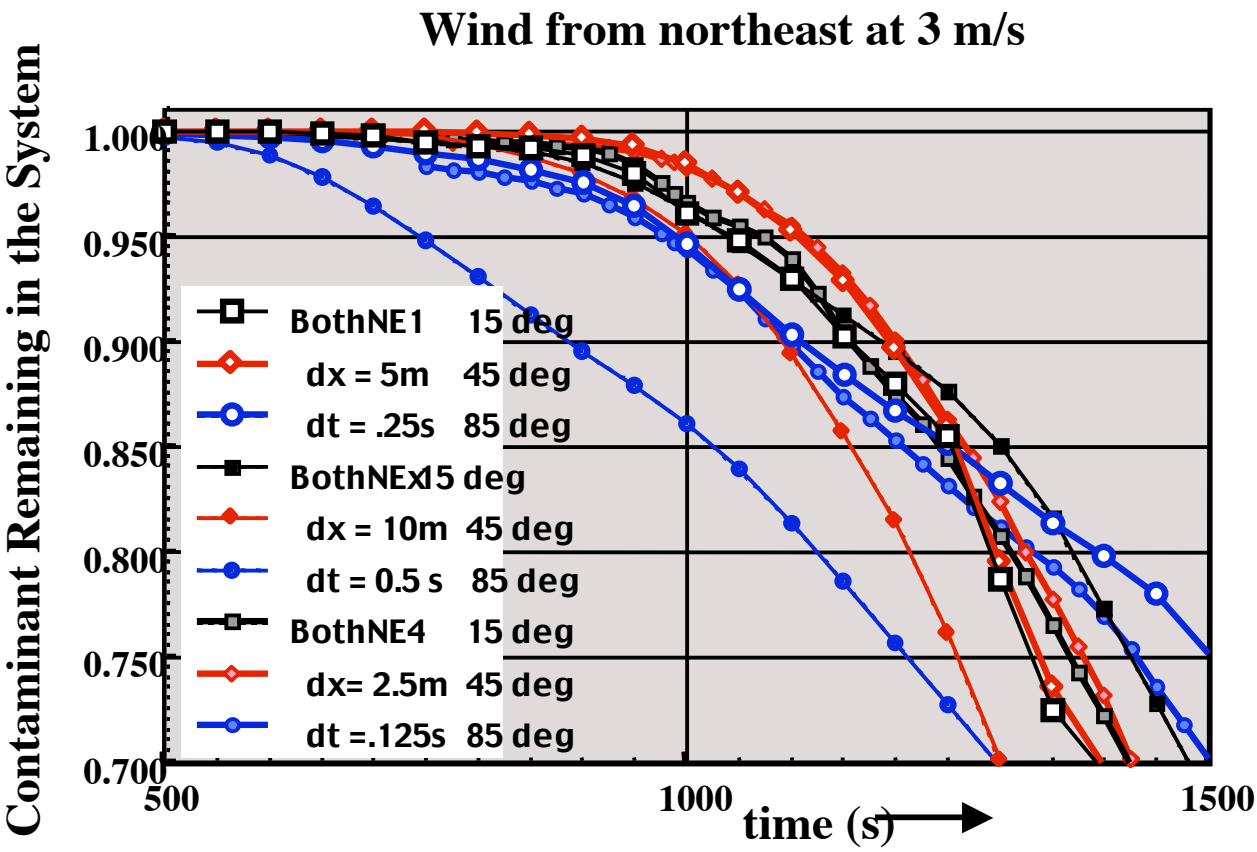

Conclusion: 5 m resolution is acceptable

Simplifications from this Con

1. Trapping of contaminant behind buildi an overall exponential decay in time

Comparison of Contaminant Decay for Different Wind Conditions

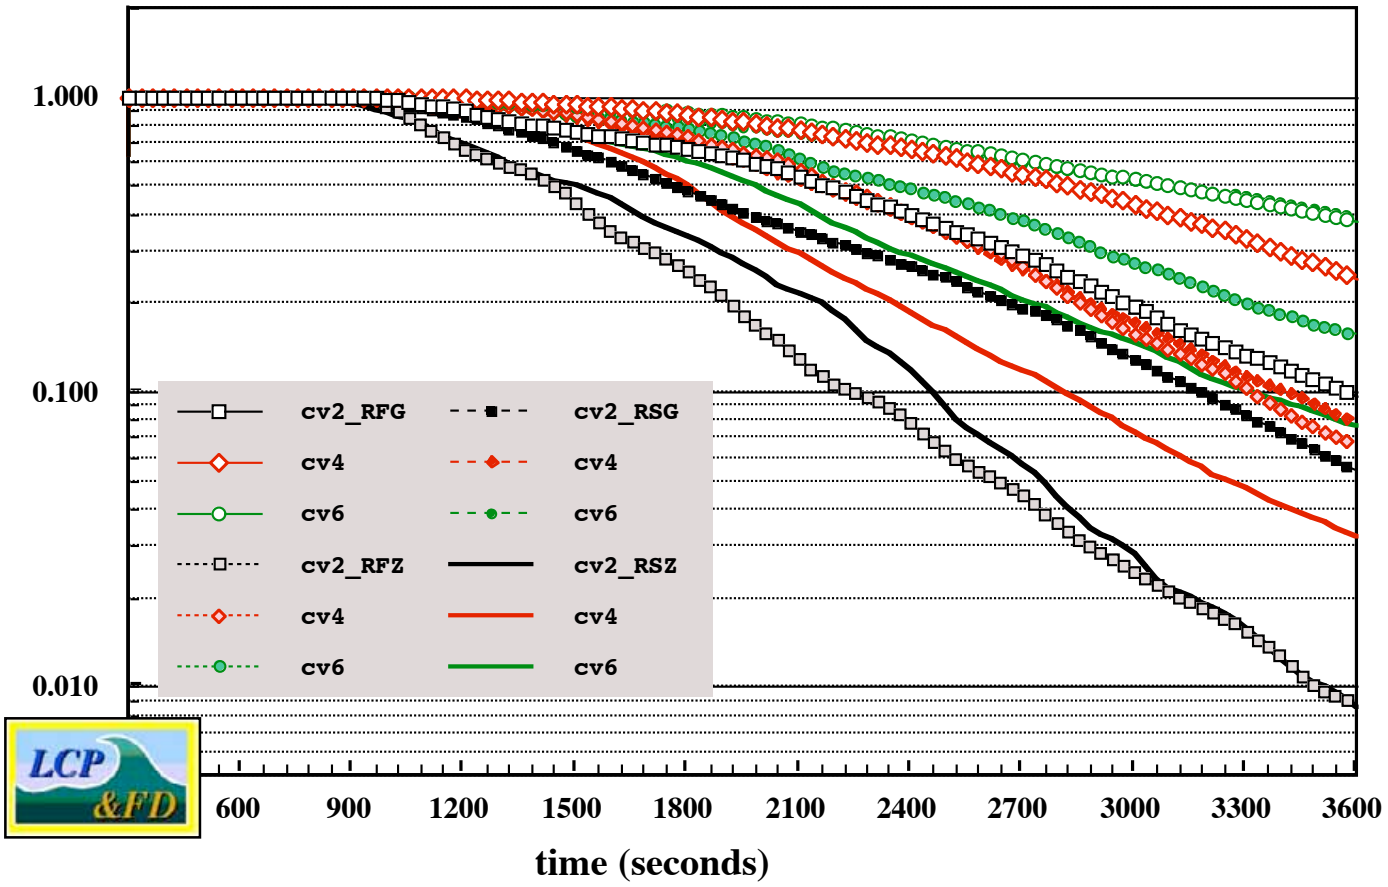

Contaminant Decay

τ_{decay} = 5 min

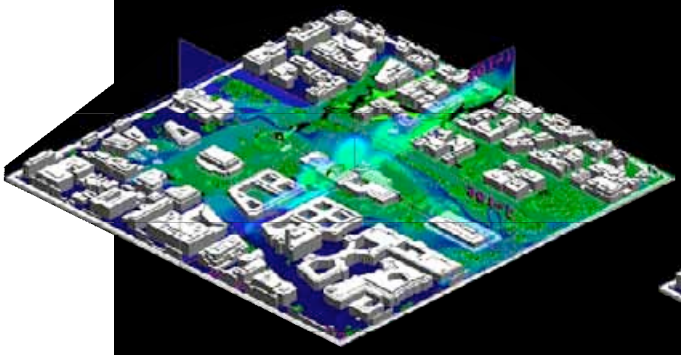

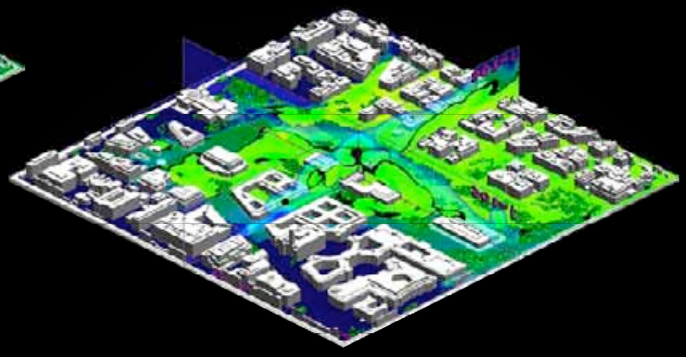

τ_{decay} = 20 min

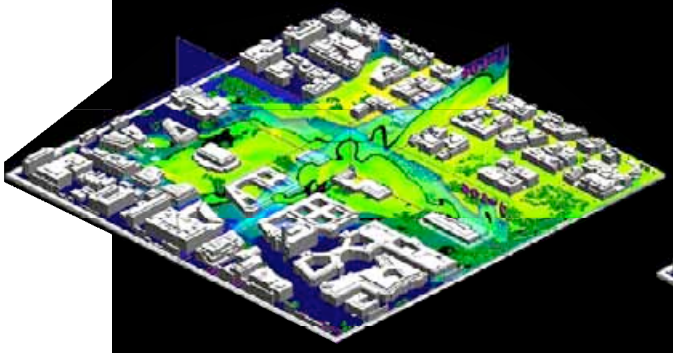

τ_{decay} = 10 min

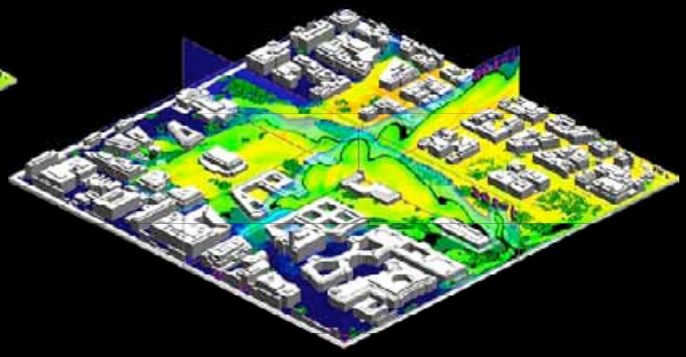

30 min after release
Wind from NE at 3 m/s
looking southeast
horizontal cut @ 11.5 m

τ_{decay} = infinity

NRL/FAST3D-CT

Contaminant Sweepout (FAST3D-CT Simulation

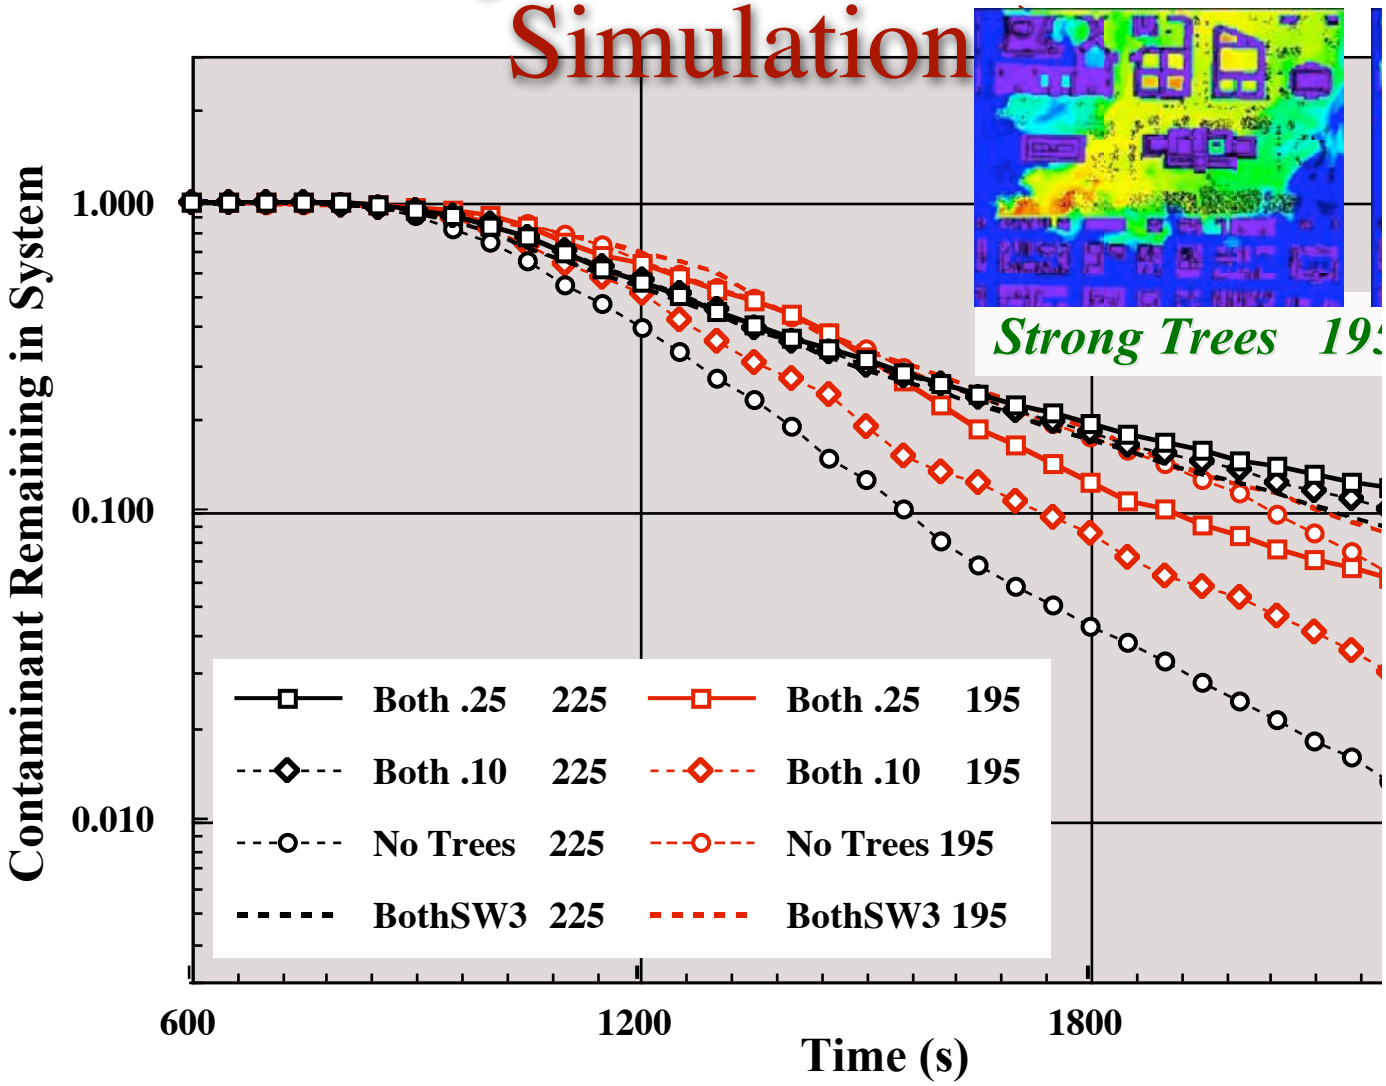

Wind from SouthWest at 3 m/s

Contaminant Decay Curves
FAST3D-CT & "Gaussian" Models

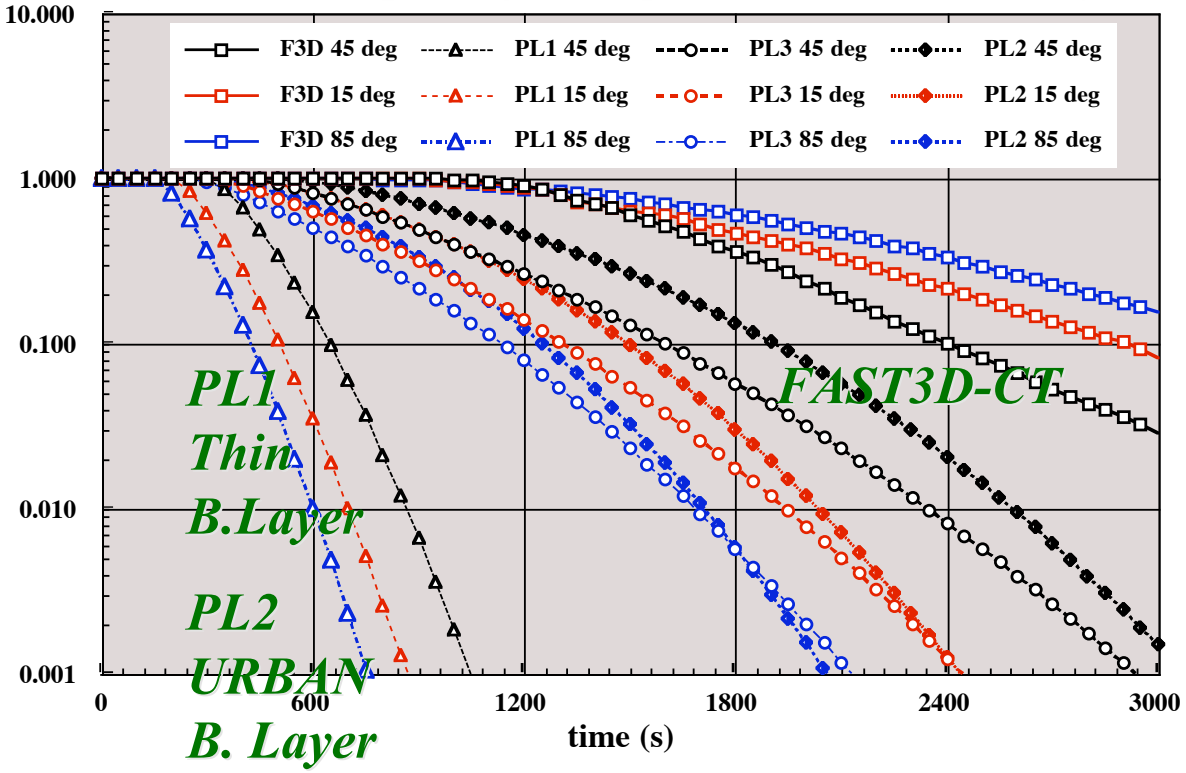

PL1 Thin B.Layer

PL2 URBAN B. Layer

PL3 Medium B. Layer

FAST3D-CT

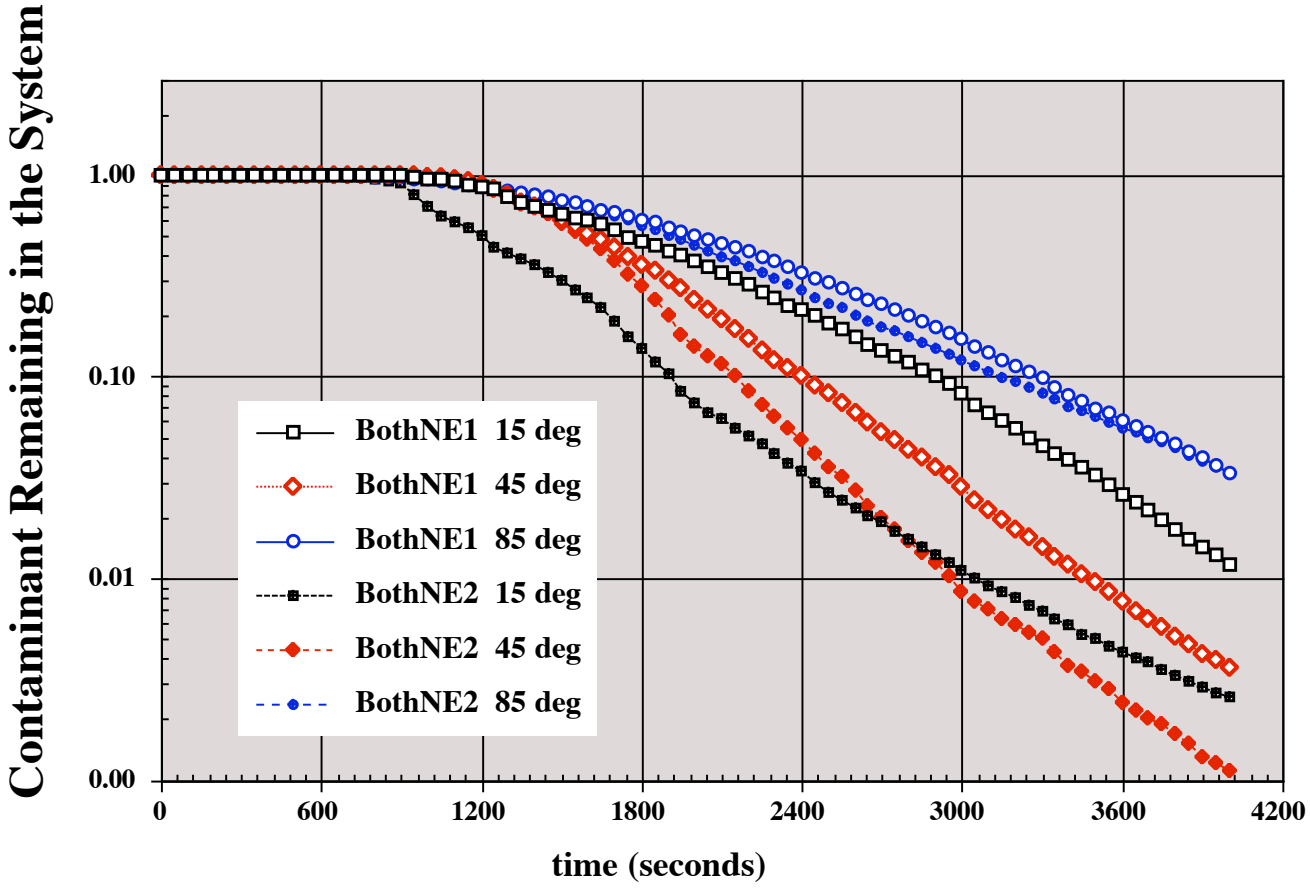

Effect of Time-Varying Winds on Contaminant Dispersion and Decay